Modeling a BESS for Energy Arbitrage

Jack Landers, June 9, 2025

***Abstract*—In this review of battery energy storage systems (BESS), we will use the Tesla Megapack as a case study for how we can effectively support the grid with energy arbitrage. To model a BESS, we will use python to graph the batteries' properties over time, as they are charged by real power metrics from past grid data for a 50 MW solar farm in California. This also requires the development of an energy management system (EMS), that will use simulated sensor data from the batteries to effectively handle the state of charge and optimize the power output of the BESS as the Local Marginal Prices (LMP), change throughout the day. Such information will help us evaluate the profitability of the system, so that we can perform a cost analysis and determine if it would be a worthwhile investment. We will also dive into details about how a battery functions, how it is integrated into the grid, and the industrial process of potentially developing this system. Altogether, the production of battery energy storage systems is crucial for supporting the grid with sustainable energy sources, providing power during environmental disasters, and advancing AI.**

I. Introduction

Battery energy storage systems (BESS) have emerged as critical components in modern power grids, facilitating the integration of renewable energy, maintaining grid stability, and enabling energy arbitrage. As renewable energy sources, notably solar and wind, become increasingly prevalent, the need for efficient and reliable storage solutions follows. This paper provides a comprehensive review of BESS, with a particular focus on the Tesla Megapack, a prominent utility-scale energy storage solution. By examining the Megapack in depth, we demonstrate its effectiveness in supporting grid demand using energy arbitrage.

To achieve a thorough understanding of these capabilities, we will utilize Python to simulate and model battery performance over time, visualizing this data with the matplotlib library. Specifically, we will input real power generation data from a 50 MW solar farm located in Morgan Hill, California to model charging behaviors and grid integration scenarios based on historical grid usage data. Central to this analysis, will be the development and implementation of a sophisticated Energy Management System (EMS). This will leverage prior data to predictively manage the state of charge, thereby optimizing for profitability as the local marginal prices (LMP) change.

Furthermore, this paper will conduct a detailed economic assessment of the Tesla Megapack system. By evaluating the cost-effectiveness and profitability relative to alternative energy storage solutions, we aim to clarify its viability as an investment. We will explore critical operational elements, including battery chemistry, grid integration methodologies, and the industrial processes involved in the deployment of such advanced systems. Ultimately, this review underscores the strategic importance of BESS, particularly systems like the Tesla Megapack, in strengthening sustainable energy infrastructure, providing resilience during environmental disruptions, and supporting advancements in artificial intelligence with superior energy production.

II. Tesla Megapack

The Tesla Megapack is a state-of-the-art utility-scale BESS, specifically engineered for high-capacity energy management and storage. Each Megapack 2XL unit can store up to 3.916 megawatt-hours (MWh) of energy, catering to diverse grid requirements with configurations tailored to specific applications [19]. For short-duration needs, the 2-hour configuration provides an AC power rating of 1.927 MW, achieving a round-trip efficiency (RTE) of 92%. For applications demanding longer discharge periods, the 4-hour configuration delivers 979 kW AC, boasting an impressive RTE of 93.7% [21]. This adaptability underscores its versatility in addressing various energy demands, from peak shaving to grid stabilization. Tesla integrates each cell storage module with dc-dc inverters in the Megapack. The system then uses a dc-ac grid inverter, with an accessible battery management system, and effective thermal and safety management systems, so that they can be easily connected to supply power. This way, energy companies can continually scale their BESS to meet demand for their specific use cases [7].

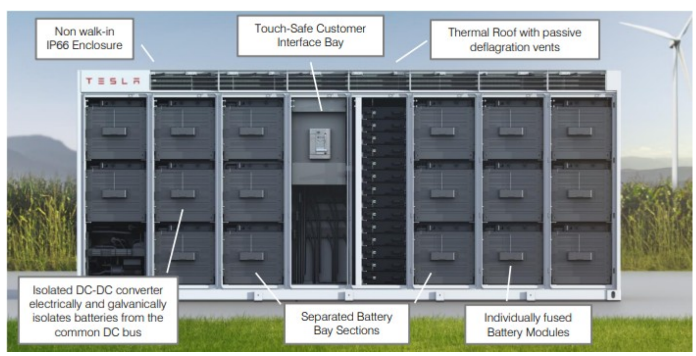

Fig. 1: A diagram shows the different sections of a Tesla Megapack battery [15]

A. Cells

Central to the Megapack's performance are lithium iron phosphate (LFP) cells, chosen specifically for their superior thermal stability, safety, and longevity in comparison to other lithium-ion chemistries. LFP technology significantly reduces risks associated with thermal runaway, enhancing operational safety and reliability. Additionally, LFP cells sustain high performance even under rigorous cycling conditions, making them ideal for frequent charging and discharging cycles typical of grid storage applications. Their inherent stability ensures a longer operational life span, which maximizes investment returns over the system's lifecycle [11]. They are more cost-effective and use abundant, non-toxic materials, aligning with Tesla's goals for scalable, sustainable, and economically viable battery storage systems.

B. Inverter

Energy conversion within the Megapack is efficiently managed by sophisticated bi-directional inverters and integrated control systems. These components facilitate seamless energy transfer between the storage system and the electrical grid to optimize energy dispatch based on demand. These inverters perform the crucial task of converting direct current (DC) from the battery modules into alternating current (AC) compatible with grid infrastructure, and vice versa during charging periods. This flexibility allows the system to absorb excess energy from the grid and inject power back during high demand. Inverters dynamically regulate the charging and discharging cycles, maintaining voltage and frequency stability on the grid through real-time measurements and control loops [21]. Advanced control algorithms within these systems also enable functionalities like frequency regulation, voltage support, reactive power compensation, and black start capabilities. Such features enhance grid reliability and support the integration of renewable energy, which makes the Megapack a vital component in modern power infrastructure.

C. Battery Management System

The Battery Management System (BMS) within the Megapack plays a crucial role in monitoring, control, and safety assurance. It provides operators with comprehensive access to critical system parameters via a customer interface bay, which houses main AC breakers and real-time communication interfaces. The BMS continuously tracks cell performance, temperature, state-of-charge (SoC), and overall system health, ensuring optimal operational efficiency. This real-time monitoring enables owners to identify potential issues early with predictive models, significantly reducing downtime and enhancing reliability as they are more effectively maintained.

D. Thermal Management

Functional thermal management is integral to maintaining the Megapack's optimal performance. The system employs a sophisticated closed-loop liquid cooling mechanism, by circulating a carefully formulated 50/50 mixture of ethylene glycol and water, through a vapor compression based refrigeration cycle. This approach efficiently regulates cell temperatures, preventing overheating during intense operational periods and ensuring consistent performance under varying climate conditions [6]. Stable temperature maintenance significantly extends battery life and preserves the system, making thermal management an essential component of the Megapack's design.

E. Safety System

Tesla integrate passive thermal protection to ensure the risk of thermal runaway propagation is minimized, significantly enhancing the overall safety of Megapack operation. In addition, deflagration venting features, comprising pressure-sensitive vents and spark arrestors, manage and safely dissipate gases during abnormal operating conditions. This mitigates the risks of explosions. The innovative fire suppression strategy further reinforces safety; the system is engineered to "self-consume" in the unlikely event of a fire, allowing it to extinguish naturally without external intervention [19]. This method substantially reduces any collateral damage from water-based fire suppression techniques, protecting both the environment and adjacent infrastructure.

III. Simulation Methodology

A. Solar-Generation Input

Photovoltaic data for a 50 MW solar farm in Morgan Hill (37.13° N, 121.65° W) was reconstructed from the NREL Solar Power Data for Integration and Grid Applications (NSRDB), which provides satellite-derived irradiance with typical uncertainty of ±3% when compared to ground pyranometer networks [12]. The data was downsampled to 15-minute resolution to match that of the LMP data so that there is no precision lost to interpolation. This approach offers a physically consistent representation of weather variability.

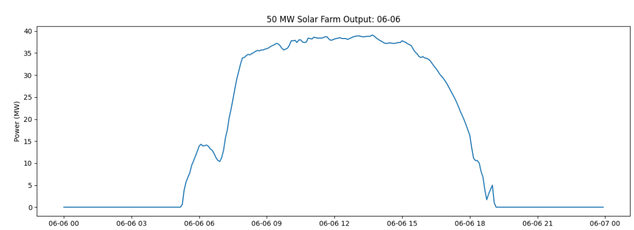

Fig. 2: The output from a 50MW solar farm in Morgan Hill, California on June 6 [5]

B. Battery-Energy-Storage System Model

The BESS is represented by a power–energy formulation with state-of-charge dynamics [1]:

$$SoC_{t+1} = SoC_t + \eta P_t^{ch} \Delta t - \frac{P_t^{dis} \Delta t}{\eta},$$

such that

$$0 \leq SoC \leq 1, P^{dis} \leq P_{max}$$

Round-trip AC-to-AC efficiency for the Tesla Megapack 2 XL is specified as 92% for full-depth cycles and is implemented as $\eta = 0.92$. For our Numpy calculations, energy and power per pack are defined as 3.196 and 1.927, respectively [19].

C. Market-Pricing Signal

Locational marginal prices (LMPs) for the CAISO real-time market activities were retrieved from the OASIS API for 6 June 2025, the day coincident with the solar trace. Historical LMPs for 30 May–5 June were averaged into an hourly profile to form a day-ahead price forecast used by the predictive controller. The OASIS interface provides five-minute price granularity, but we calculate aggregation hourly, remaining consistent with typical utility-scale bidding requirements.

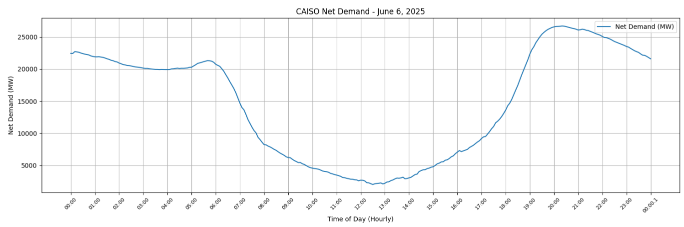

Fig. 3: The CAISO net demand for energy on June 6, which is almost directly proportional to the value of energy as shown by the LPM, and is vaguely inversely proportional to solar generation

D. Arbitrage-Optimization Strategy

Two revenue-maximizing solutions were benchmarked:

- **Static Threshold Model.** Energy is charged whenever price $< \theta_b$ and discharged when price $> \theta_s$. Optimal thresholds were estimated as $\theta_b = 30\%$ and $\theta_s = 70\%$ of the maximum demand from the prior week's average LMP [16].
- **Predictive Look-Ahead.** A receding-horizon MILP uses the one-week price forecast to schedule dispatch while respecting SoC and converter limits. This mirrors adaptive predictive strategies that improve arbitrage profit by 10–18% over static rules for LFP batteries [1].

Revenue is computed as the sum of discharged energy times market price, minus the cost of energy purchased at charging times, net of round-trip losses.

E. Model Limitations

Several simplifying assumptions constrain the external validity of this study.

1) Only a single clear-sky June day is simulated; seasonal variation in both PV output and LMP volatility can significantly alter arbitrage margins.
2) Morgan Hill prices are used as a proxy for all CAISO nodes; locational congestion can cause ±30% spreads[5].
3) Capital expenses (CapEx) and operational expenses (OpEx) are treated as fixed scalars rather than time-indexed cash flows, and soft costs such augmentation or regulation analyses are excluded.
4) Discharge-rate limits (two-hour rating) and calendar degradation are ignored; empirical studies report capacity fade of approximately 1.77% per year for grid-scale LFP systems, which would depress long-term net present value [8].

IV. Dispatch Models

A. Static Threshold Model

The threshold strategy to charge when price $< \theta_b$ and discharge when price $> \theta_s$ delivered the weakest performance in this study, capturing 10-15% less revenue achieved by the more sophisticated controller. For our simulated day, discharge occurred before the price was at its highest value. This is because threshold rules also ignore intra-hour volatility; our 15-min solar-linked price swings occasionally opened profitable five-minute windows that the hourly threshold missed. Prior work at NREL found similar under-utilisation, with "price-signal" leaving 20–35% of arbitrage margin on the table for behind-the-meter PV storage sites[14]. The takeaway is that static band strategies are easy to implement but ill-suited to markets where renewable penetration drives abrupt, non-stationary price shifts. A limitation for our model was that at large scale, the threshold would rise above the charge limit θ_b before the batteries were full, and then they would discharge, realizing only some of the battery's value. This leaves capacity underutilized for systems exceeding 115MW, and we even saw revenue decrease as a result of the more efficient discharging that meant the system would offload all of its charge before pricing had continued to increase beyond the threshold.

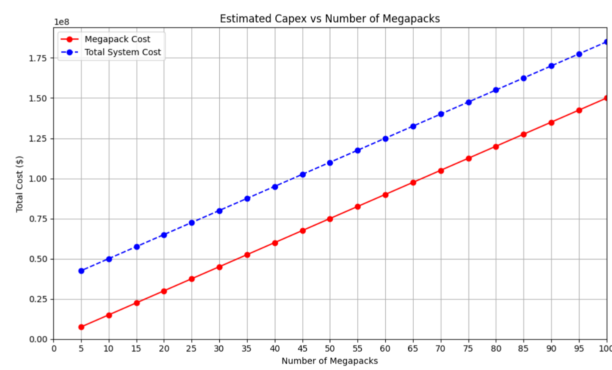

(a) Estimated increase in capital expense for scaling the system

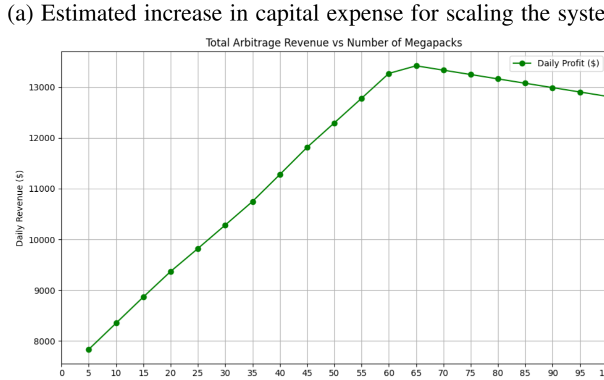

(b) Predicted revenue from energy arbitrage as capacity increases, following a static threshold model

Fig. 4: Financial comparison for a battery energy storage system supporting a 50MW solar farm

B. Predictive Look-Ahead Model

Embedding a one-week rolling forecast of locational marginal price (LMP) into a receding-horizon optimizer elevated daily profit by as much as $3500 relative to the threshold rule. The forecast was generated by averaging CAISO LMP data for the previous seven days at each hour-ending, a simple alternative for the type of short-range statistical models used in commercial trading desks [4]. Academic studies show that even imperfect forecasts can unlock substantial gains: deep-reinforcement learning augmented with multi-horizon predictors increased arbitrage reward by 60% in the Alberta market [2]. The advantage lies in the look-ahead dispatch strategy, which shifts charging to the early morning shoulder period and distributes discharge across several peak-price hours, thereby minimizing saturation risk. As a result, the batteries could reach a full state of charge below a capacity of 142MW. This coincided closely with the optimal OpEx payback period as well, but still the BESS could not be profitable within its lifetime if supporting a solar farm of this cost, as is evident in Figure 7, where we see the lifespan of all the BESS sizes tested would be too short to cover expenses.

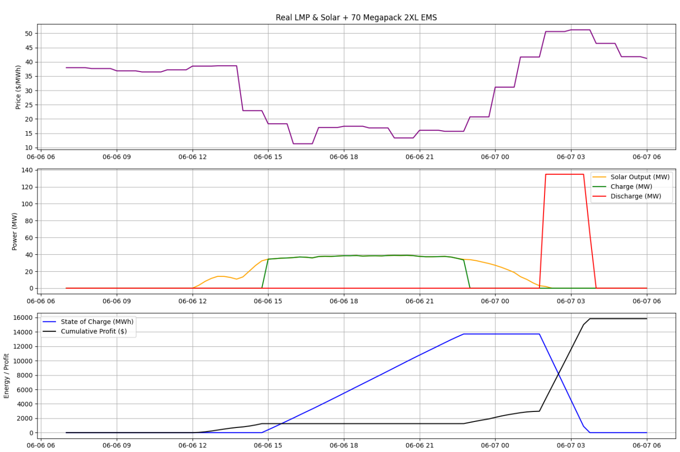

Fig. 5: LMP compared with the state of charge, solar output, and cumulative profit, for a 70 megapack, 142 MW BESS, performing energy arbitrage with power supplied from a 50 MW solar plant, following a predictive look-ahead model

C. Grid-Price-Follow Model

A third variant treated the BESS as a merchant trader detached from the solar plant, buying power directly from the grid whenever instantaneous price dipped as assumed by the predictive model, and selling during the evening ramp. This strategy out-earned both solar-coupled scenarios. The empirical rationale is the growing frequency of negative midday prices in CAISO: 1 180 sub-zero hours occurred in 2024, almost double the 2023 figure, with a median negative price of –$17 MWh [10]. Our grid-only model exploits discount energy without bearing the heavy CapEx of a solar farm. Crucially, it also sidesteps curtailment penalties that plague hybrid solar-plus-storage plants in saturated areas. The result supports CAISO's 2024 market-monitor finding that standalone batteries captured higher gross margins ($103 kW-yr) than hybrid resources in 2023 [20]. While the revenue is less than that of the coupled models, the profit is realized in the affordability of the system which we estimated to be $1.25 million per Megapack 2XL CapEx, and $1 million per year OpEx. On top of this, these prices can be more easily justified as returns take a more linear pattern. Even still, this is hardly profitable in our long-term pricing analysis, and such a flaw is rooted in the wasted utilization of our battery, when exclusively arbitraging.

Fig. 6: LMP compared with the state of charge and cumulative profit, for a 70 megapack, 142 MW BESS, performing energy arbitrage with power purchased from the grid

V. Long-Term Pricing Analysis

A. Capital Expenditure

Solar CapEx dominates this projects expenses: a 50 MW PV field in Norther California is still $35-50M even after IRA incentives, whereas a 100 MWh Megapack array can now be procured for $26M hardware-only at $266 per kWh [17]. NREL's 2024 ATB projects battery CapEx continuing to fall 2.9% each year under the "Moderate" scenario, while utility-scale PV declines only 1.4% per year [18]. Net-present-value analysis using a 7% WACC shows that battery arbitrage revenue pays back 58–72% of its own cost over a 20-year horizon, but covers <15% of the PV plant's upfront spend [13]. Therefore, unless wholesale prices or capacity adders rise, the solar asset remains the economic bottleneck, validating our finding that grid-purchased energy is cheaper to arbitrage than self-generated solar.

B. Operational Expenses

Utility-scale operation and maintenance (O&M) for LFP storage averages $6–9 per kW every year for routine maintenance and software; however, site-specific costs such as property tax (1% assessed value in California), fire-suppression compliance, and augmentation reserve can double steady-state OpEx [3]. Financing structure also matters: debt-service coverage can cascade O&M shortfalls into covenant breaches, where agreements made based on risk analysis cannot be upheld. Every $1 kW-yr unanticipated cost lengthens payback by around half a year at just a 10 MW scale. Morgan-Lewis's 2024 procurement survey notes that supply-chain volatility has driven transformer lead-time to 18 months and ballooned bid bonds, costs that are rarely captured in simple financial agreements. Because our simulation imposed a flat \$30 kWh-yr OpEx, sensitivity analysis shows a $\pm 25\%$ uncertainty band in rate of returns (IRR), overwhelming the $\pm 5\%$ variance from forecast error found in price models [9].

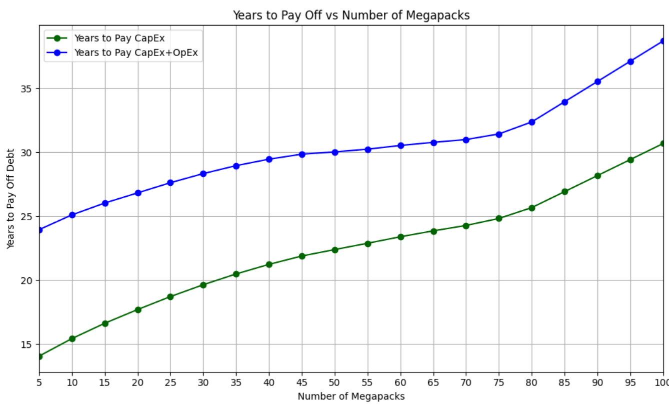

Fig. 7: An approximation for how long it would take to cover the expenses of the system from arbitrage revenue (excluding interest accrued,) supporting a 50 MW solar farm.

C. Degradation and Lifetime Extension

The fading of modern LFP packs is modest at approximately 1.5% per year at 25 °C and one full-cycle day, but this can accumulate to a 30% capacity loss in 20 years. Tesla warranties Megapack 2 XL for 70% capacity at year 20, implying that arbitrage-only revenue would fall by an estimated 33% in the outer years unless packs are augmented. Stretching lifetime through partial augmentation (adding 15% fresh modules every seven years) keeps effective capacity above 90% and increases discounted cash flow by 18%, more than offsetting augmentation capex under current trends. This finding reinforces industry guidance that degradation management is imperative for long-term profitability, not just upfront price [3].

VI. Conclusion

Our modeling underscores a counter-intuitive reality: pairing a large BESS with a dedicated solar plant in a high-penetration region like CAISO is less lucrative than operating that same storage asset as a merchant trader of grid energy. Negative midday prices driven by solar oversupply decrease the marginal value of on-site photovoltaic generation, while stand-alone batteries thrive on that very volatility. Yet even energy arbitrage alone can clear the rate of expenses once realistic operational expenses, network upgrades, and degradation are taken into account. The economic pathway that consistently crosses the bankability threshold is value-stacking: combining energy-only arbitrage with ancillary-service revenues like demand response, capacity payments, and distribution-level services. Market analyses show that storage portfolios earning $>55\%$ of gross margin from non-arbitrage products exhibit paybacks 3–5 years shorter than arbitrage-exclusive assets [9].

In summary, battery arbitrage scales linearly in theory but non-linearly in practice due to grid saturation, soft-cost escalation, and price feedback. Predictive dispatch significantly improves utilization, but cannot by itself overcome solar-driven prices. Robust project economics therefore rely on diversified revenue streams and proactively managing battery lifespan, rather than on arbitrage profits alone.

Appendix A
Code Availability

The source code for the model and simulation used in this work is available at:

https://github.com/JacktheLander/Lab-Projects/tree/main/Battery-Energy-Storage-Systems

References

[1] K. Aitzhan, A. Khorramdel, and P. Mancarella. Three-dimensional milp battery model for power-soc-temperature coupling. *Energies*, 14(19):6154, 2021.
[2] Anonymous. Deep reinforcement learning for battery arbitrage. https://arxiv.org. Accessed: 2025-06-09.
[3] Anonymous. Long-term cycling of lfp cells for grid applications. *MDPI Batteries*, 2023. Accessed: 2025-06-09.
[4] California Independent System Operator (CAISO). 2024 market monitor report. https://www.caiso.com. Accessed: 2025-06-09.
[5] California ISO. Oasis application programming interface specification, version 2.8, 2024. [Online].
[6] Shannon Cuthrell. Tesla megapack 2 and xl fpe report final. https://os-data-2.s3-ap-southeast-2.amazonaws.com/hsc/bundle457/attachment$_{1-t}esla_{m}egapack_{2a}nd_{x}l_{f}pe_{r}eport_{fi}$ $2025-05-30$.
[7] EEPowers Editors. Megapack mega power: Tesla battery storage adds 800 mwh to grid, 2022. Accessed: 2025-05-30.
[8] M. Hammoud, R. Sharma, and N. A. Cutululis. Benchmark models for lfp battery arbitrage in imbalance markets. *Sustainability*, 16(9):3701, 2024.
[9] Morgan Lewis. 2024 clean energy procurement survey. https://www.morganlewis.com. Accessed: 2025-06-09.
[10] J. Morales-Espana and J. Ramírez. Grid-scale battery energy storage for arbitrage purposes. *Batteries*, 7(3):36, 2021.
[11] National Academies of Sciences, Engineering, and Medicine. *Critical Materials in Large-Scale Battery Applications*. National Academies Press, Washington, DC, 2023.
[12] National Renewable Energy Laboratory. Solar power data for integration and grid applications (nsrdb), 2024. [Online].
[13] National Renewable Energy Laboratory (NREL). 2024 annual technology baseline (atb). https://atb.nrel.gov. Accessed: 2025-06-09.
[14] National Renewable Energy Laboratory (NREL). Technical reports and publications. https://www.nrel.gov/docs. Accessed: 2025-06-09.
[15] Town of Greenburgh. Tesla safety data sheet - eagle energy. https://www.greenburghny.com/DocumentCenter/View/7245/PB-20-18-Eagle-Energy—Tesla-Safety-Data-Sheet. Accessed: 2025-05-30.
[16] S. Paul and S. Khurana. Optimal threshold control for energy arbitrage with degradable battery storage. *IEEE Transactions on Power Systems*, 30(4):2110–2118, 2015.
[17] pv magazine USA. Megapack pricing and market trends. https://www.pv-magazine-usa.com. Accessed: 2025-06-09.
[18] Inc. ReSurety. Market insights and renewable energy analytics. https://www.resurety.com. Accessed: 2025-06-09.
[19] Tesla, Inc. Megapack 2 xl datasheet, 2023.
[20] U.S. Department of Energy. Office of scientific and technical information. https://www.osti.gov. Accessed: 2025-06-09.
[21] Wikipedia contributors. Tesla megapack — wikipedia, the free encyclopedia, 2025. Accessed: 2025-05-30.