Optimization of Microgrid Design for Data Centers

Max Bryggman, Ella Christensen

Master's thesis
2025:E21

Lund University

Faculty of Engineering
Centre for Mathematical Sciences
Mathematical Statistics

Centrum Scientiarum Mathematicarum

Master's Theses in Mathematical Sciences 2025:E21
ISSN 1404-6342

LUTFMS-3515-2025

Mathematical Statistics
Centre for Mathematical Sciences
Lund University
Box 118, SE-221 00 Lund, Sweden

http://www.maths.lu.se/

Abstract

The rapid adoption of AI technologies has significantly increased the global demand for data centers. Data centers require a stable and reliable power supply. As the pressure on national electrical grids intensifies, data centers may need to seek alternative solutions to cover their power demand. One such solution is through integration of microgrids — localized energy systems capable of operating independently or connected to the main grid.

This thesis explores the configuration of a microgrid consisting of wind power, solar (photovoltaic) power, and batteries. A combination of multiple optimization techniques; Particle Swarm Optimization, the Simplex method, and Branch and Bound, is applied to historical data to determine the cost-optimal configuration of the microgrid. In addition, Monte Carlo simulations with Extended Moving Block Bootstrapping are conducted to generate scenarios for future solar and wind power production.

The simulations are used to evaluate the robustness of the optimized solution under varying weather conditions and to perform a sensitivity analysis of the microgrid's performance. The results indicate that, under the assumptions made in this report, purchasing the entire demand from the spot market is more economically favourable than investing in a microgrid. However, if full reliance on the spot market is not possible, the most optimal microgrid configuration should primarily consist of wind power, complemented by a smaller share of solar power and battery storage.

Keywords: Microgrid, Particle Swarm Optimization (PSO), Bootstrap, Wind power, Photovoltaic power

Acknowledgements

We would like to express our gratitude to our supervisor, Erik Lindström, for his valuable guidance, insightful feedback, and continuous support throughout this thesis.

We also extend our thanks to Bodecker Partners for providing us with insight into the energy market, relevant data, and thoughtful input. A special thanks goes to Johan Claesson and Niklas Gälldin for their stimulating discussions and constructive suggestions, which greatly contributed to the development of this work.

Contents

1 Introduction

1.1 Background

The rapid growth of modern society and the rapid adaptation of AI usage have significantly increased the demand for data centers. McKinsey and Company's analysis suggests that the global demand for data center capacity is expected to increase annually at a rate of 19 to 22 percent between 2023 and 2030 (Srivathsan et al., 2024). As more data centers are established, the power consumption increases accordingly. This growing demand puts pressure on the electricity grids, and in combination with the environmental targets, it may become an obstacle for new data centers.

Data centers require a constant and reliable power supply to maintain operations, known as their baseload. This baseload can be secured through Power Purchase Agreements (PPAs) or by purchasing electricity from the grid. However, because of the already high, and constantly growing, pressure on the electricity grid, these sources may not always be sufficient. In such cases, data centers may use microgrids, small localized energy systems that integrate multiple energy sources to ensure a reliable energy supply. Microgrids can combine renewable energy sources, such as wind and solar power, with energy storage solutions, like batteries, to enhance energy security and sustainability.

1.2 Objective

The objective of this thesis is to design a microgrid that integrates wind power, solar power, and battery storage to provide a data center with its baseload electricity demand. The focus is on identifying the most cost-optimal configuration for the microgrid. The study is specifically limited to a microgrid operating one year ahead within the bidding zone SE2 (see Figure 1).

Figure 1: The four electricity bidding zones in Sweden (Trustio AB, 2025).

2 Theory

2.1 Microgrid

A microgrid is a localized energy-producing system that operates independently of the main power grid. They can be either completely isolated, then often referred to as energy islands, or they can have a connection point to the main power grid. Microgrids combine multiple energy sources, often wind and solar power, with energy storage solutions like batteries. This enables facilities to generate sustainable power, ensure a more reliable energy supply and, optimize energy costs by storing electricity and selling excess power back to the grid when spot prices are high.

2.1.1 Wind Power

Wind power is a renewable energy source that generates electricity through the energy of the wind. Wind power production is seasonal with higher production in winter months and during nighttime. In 2024, wind power contributed 24% of Sweden's total electricity production. The majority of Sweden's wind power generation comes from SE2, which accounted for 42% of the country's total wind power output in 2023 (Energimyndigheten, 2024).

2.1.2 Solar Power (Photovoltaic Power)

Solar energy can be converted into electricity either through thermal methods or by using photovoltaic (PV) cells. In the context of microgrids, photovoltaic technology is the most common method. Consequently, this thesis will refer specifically to electricity generated from sunlight by the microgrid, as photovoltaic power.

In 2024, solar power accounted for 2.4% of Sweden's total electricity production (Svensk Solenergi, 2025). However, its generation is highly dependent on seasonal variations, with significantly lower output during the winter months in SE2.

2.1.3 Batteries

One of the main challenges of both wind and photovoltaic power is their *intermittency*, meaning production cannot be guaranteed at all times, (Ekoue et al., 2025). Photovoltaic power generation drops significantly during winter and ceases completely at night. Even during the day, cloud cover can reduce photovoltaic output. Similarly, wind power generation depends on wind conditions, which can be unpredictable.

To address these challenges, batteries play an important role in storing excess energy when production is high and releasing it when demand exceeds supply. This ability to balance energy supply and demand has driven rapid growth in the battery storage market. According to Jarbratt et al. (2023), the global market for Battery Energy Storage Systems (BESS) is expected to double between 2023 and 2030. By storing surplus renewable energy, batteries help reduce reliance on non-renewable energy sources, contributing to a more sustainable and resilient energy system.

In addition to improving system reliability, batteries also offer a way to increase profitability through energy arbitrage. Since renewable energy generation often coincides with lower electricity prices, and prices tend to be higher during periods of low renewable output, battery systems can buy or store electricity when prices are low and sell or discharge when prices are high. This price arbitrage potential could make batteries an attractive component in a microgrid that uses renewable energy sources (Gaspar et al., 2021).

There are currently multiple types of batteries on the market. In the study "Optimisation and economic feasibility of Battery Energy Storage Systems in electricity markets: The Iberian market case study" by Gaspar et al. (2021), an economic viability evaluation was conducted for different battery types. The results showed that only lithium-ion batteries were considered a profitable investment. Therefore, lithium-ion is the battery type assumed in this report.

2.2 Electricity Market

The physical electricity market in Sweden can be divided into four sub-markets that have four different time frames (Svenska Kraftnät, 2025). The first one is the PPA (power purchase agreement) market, or the price-securing market. This is often done long before the actual delivery of the electricity and is a way for energy suppliers to secure energy prices early in advance. The second market is the **spot market**, also known as the **day-ahead market**. Trading on the day-ahead market is, as the name suggests, done one day before the delivery of the electricity. At 11 am, the market participants place an order for all of the hours the following day. The suppliers estimate how much they can sell and the buyers how much they demand. From this the spot prices are set to the equilibrium prices and published for all hours the next day.

The electricity price for each hour is determined by the intersection of the supply and demand curves. In the energy market, the supply curve has a stepwise structure, resembling a staircase, see Figure 2. The hourly market price is set by the marginal bid of producing the last kilowatt-hour needed to meet the demand (Energimarknadsinspektionen, 2023).

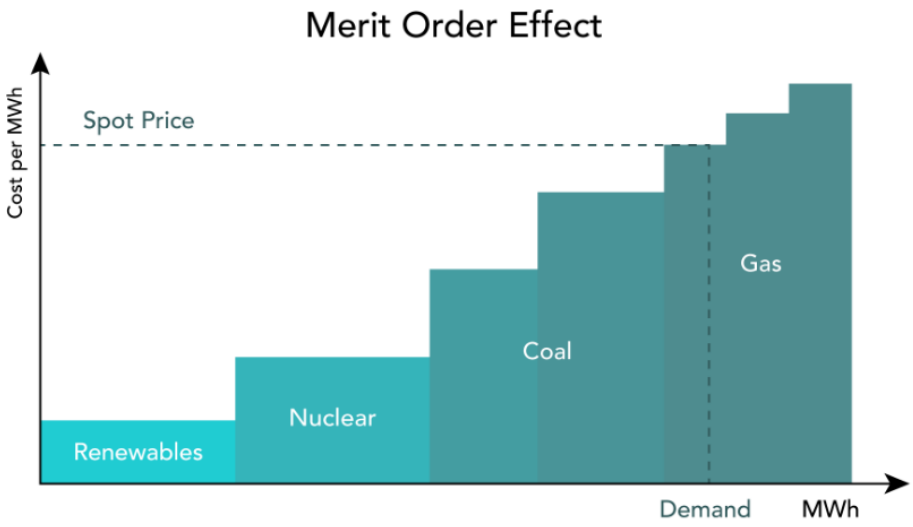

Figure 2: The Merit Order Effect, illustrating spot price determination by the price bid by the marginal generator (Bowden, 2023).

Under this pricing model, renewable energy sources are utilized first, as their low marginal costs make them the most economically favorable option. However, when the supply from renewables is insufficient, either due to low production or high demand, additional production is required. In such cases, nuclear power is brought online, which typically has a higher cost, thereby raising the market price. If demand still exceeds supply, fossil fuel-based generation is activated, further increasing the price due to its even higher marginal cost.

In some situations, the market dynamics can even result in negative electricity prices. This typically occurs when there is low demand combined with high production from renewable sources like wind and solar. Since electricity must always be balanced in real-time and is difficult to store, producers might choose to sell electricity at a negative price instead of shutting down their production. These situations highlight the need for flexible energy systems that can respond quickly to changes in supply and demand.

However, since the spot market is built on forecasts and it is hard to predict the future, there are two further markets. The first one is the **intra-day market** where electricity market participants can buy and sell electricity up to one hour before delivery (Lindholm, 2023). Lastly, to make sure that all produced electricity is consumed, there is a **balance market**, which in Sweden is run by Svenska Kraftnät. This market is active before and during the time of delivery of the electricity.

Although the electricity market offers various opportunities for optimizing production, this thesis will focus solely on the day-ahead market.

2.3 Grid Limitations

Sweden's electricity grid is divided into three main components: the transmission grid, the distribution grid, and interconnections with other countries. The transmission grid transports large amounts of electricity from major power producers to regional distribution networks. It spans the entire country and can be compared to highways for electricity, delivering power to smaller regional and local roads, represented by the distribution grid (Svenska Kraftnät, 2024b).

The trade of electricity is limited by how much power safely can be transferred within and between these networks. In Sweden, the responsibility for these calculations lies with Svenska Kraftnät, which ensures that transmission lines can handle the electricity flow without exceeding their capacity (Svenska Kraftnät, 2024a). The amount of electricity a producer can import from or export to the grid is determined by its point of connection (POC). The specific capacity of this connection depends on various factors, such as local grid conditions, available transmission capacity, and other requirements set by the grid operator.

A microgrid can either operate independently, disconnected from the main grid, or it can operate with a grid connection. In cases where a microgrid is grid-connected, the limitations imposed by the point of connection play a significant role. These restrictions determine how much electricity the microgrid can trade with the main grid. A limited import capacity may increase the need for battery storage to ensure a sufficient energy supply during periods of low production. On the other hand, a restricted export capacity can limit the ability to sell surplus energy from photovoltaic and wind power. This could lead to curtailments (active reductions in production) or penalties for exceeding export limits.

2.4 Simulation

Monte Carlo simulations are a method used to estimate functions and model the behaviour of complex systems (Harrison, 2010). This is done through random sampling. Samples can be drawn from various probability distributions, such as Normal, Uniform, Exponential, or Poisson distributions, or generated through bootstrapping from previous observations. There are multiple ways to apply Monte Carlo simulations, one of which is Moving Block Bootstrapping (MBB).

2.4.1 Moving Block Bootstrapping (MBB)

A proposed method for wind power simulation is the Moving Block Bootstrapping (MBB) method (Usaola, 2014). MBB is a Monte Carlo-based resampling technique that generates new time series by sampling overlapping blocks of historical data. This method is particularly useful for preserving autocorrelation and other time dependent structures in data.

Unlike parametric simulation models, MBB does not assume an underlying probability distribution, making it purely data-driven and free from distributional assumptions. A key limitation of MBB is that it cannot generate more extreme events beyond those already

observed in the historical dataset. This means that if the method should yield statistically significant results, a sufficiently long time series is required. If the historical data set is too short, the method may underestimate variability and fail to capture rarer events.

In the article *Synthesis of hourly wind power series using the Moving Block Bootstrap method* by Usaola (2014), the MBB method is applied to simulate wind power data that has been normalized with installed capacity and, the seasonality pattern has been removed. Blocks of multiple consecutive days are then sampled from the same season of historical data to ensure consistency in seasonal wind behaviour. This modification of the MBB method with sampling from the same season is referred to as EMBB (Extended Moving Block Bootstrap)(Dudek, 2018). This builds on the MBB method by addressing one of its key limitations, the inability to preserve and handle seasonal structures in time series data. While MBB works well for non-stationary data, it doesn't explicitly account for periodic patterns, which can be crucial in datasets like wind power, where seasonal variations play a significant role. EMBB extends the original MBB by ensuring that resampled blocks of data are sampled from the same season, maintaining the seasonal characteristics of the data.

This extension makes EMBB particularly suitable for applications with strong seasonal components, as it preserves the periodicity within the simulated time series. Additionally, EMBB allows for better handling of complex time series that exhibit both long-term trends and seasonality, ensuring that these crucial aspects of the data are preserved in the simulations (Dudek, 2018).

One important aspect of block bootstrap methods is the choice of block size. The size of the blocks influences the results of the simulations significantly. If the blocks are too short, the simulated time series may fail to capture long-term dependencies and variability. Conversely, if the blocks are too long, the generated series may become overly similar to the original data. This will reduce the randomness that is needed for reliable simulations.

The study *Optimal choice of bootstrap block length for periodically correlated time series* by Bertail and Dudek (2024), examines different choices of block lengths, the optimal block length to minimize the Mean Squared Error (MSE) of the block bootstrap samples is found to be proportional to $n^{\frac{1}{3}}$ for the EMBB method. This recommendation can be used, but it is important to note that there is no universally optimal block length. The optimal block length may vary for different problems.

2.4.2 Stationarity of a Timeseries

Statistical modelling based on historical data using methods like AR, ARMA, or SARIMA models is an alternative to Monte Carlo-based simulation. These models allow for simulating future values by capturing the underlying structure of the data. However, such approaches typically assume that the time series is stationary, as non-stationary data can lead to unstable or diverging simulations.

A time series is considered stationary if its statistical properties remain constant over time. More precisely, a process must fulfil the following three conditions to be classified as stationary (Jakobsson, 2021):

(i) The mean function is constant and the constant value is finite, i.e., $m_x(t) = m_x < \infty$.

(ii) The covariance $\text{Cov}(x_t, x_s)$ is only dependent on the difference $(s - t)$.

(iii) The variance of the process is constant and finite, i.e., $\mathbb{V}[x_t] < \infty$.

To assess whether a time series satisfies the stationarity assumption, several statistical tests can be employed.

2.4.2.1 ARCH-test To test whether a time series exhibits constant variance, the ARCH test (Autoregressive Conditional Heteroscedasticity) can be applied. The test was introduced by Engle (1982), who was awarded the Nobel Prize in Economics for his contributions to the analysis of time series with time-varying volatility in 2003.

The ARCH test is designed to detect heteroscedasticity, which means that the variance of the residuals is not constant over time. In contrast to homoscedasticity, where the variance remains constant, heteroscedasticity implies that certain periods exhibit higher volatility than others.

The ARCH test is applied through the following steps:

1. Consider a time series x_t, which could be raw data or residuals from a fitted model.
2. Square the series to obtain a new series:

$$u_t = x_t^2$$

3. Regress u_t on its own lagged values. For example, an ARCH(1) model is written as:

$$u_t = \gamma_0 + \gamma_1 u_{t-1} + \nu_t$$

and a general ARCH(q) model is written as:

$$u_t = \gamma_0 + \gamma_1 u_{t-1} + \gamma_2 u_{t-2} + \cdots + \gamma_q u_{t-q} + \nu_t$$

In this context:

- γ_0 is the intercept term,
- γ_i for $i = 1, \ldots, q$ are the coefficients on the lagged squared terms,
- ν_t is the error term of the regression.

The hypotheses of the ARCH test are:

$H_0 : \gamma_1 = \gamma_2 = \cdots = \gamma_q = 0$ (There are no ARCH effects/homoskedasticity in the data)
$H_1 :$ At least one $\gamma_i \neq 0$ (Presence of ARCH effects/heteroskedasticity in the data)

To perform the test, compute the following test statistic:

$$\text{LM} = T \cdot R^2$$

where:

- T is the number of observations,
- R^2 indicates how much of the variance in u_t is explained by its own lagged values in the regression.

Under the null hypothesis, the LM statistic follows a chi-squared distribution:

$$\text{LM} \sim \chi^2(q)$$

If the test statistic exceeds the critical value of the chi-squared distribution with q degrees of freedom, then H_0 is rejected, indicating the presence of ARCH effects in the data.

The presence of ARCH effects violates the assumption of constant variance, which is required for certain models and stationarity conditions.

2.5 Optimization

Optimization is a mathematical concept about finding the best possible solution from a set of different feasible solutions. There are numerous approaches to optimization. There are exact methods that provide optimal solutions, but there are also heuristic methods that provide solutions close to the optimal solution, but at a lower computational cost. In this section, three optimization methods are presented: The simplex method and the Branch and bound method for solving mixed-integer linear problems (MILP) and Particle Swarm Optimization (PSO), a heuristic algorithm inspired by swarm intelligence.

2.5.1 Solving MILP

There are multiple ways to solve a linear optimization problem, but when the problem involves binary variables, it becomes a Mixed-Integer Linear Programming (MILP) problem.

Definition of a Mixed-Integer Linear Programming (MILP) Problem:

$$\begin{aligned} \text{Minimize} \quad & c^T x \\ \text{subject to} \quad & Ax \leq b \\ & x_i \in \{0, 1\}, \quad i \in \mathcal{I} \\ & x_j \in \mathbb{R}, \quad j \notin \mathcal{I} \end{aligned}$$

Where:

- $x \in \mathbb{R}^n$ is the vector of decision variables.
- $\mathcal{I} \subseteq \{1, 2, \ldots, n\}$ is the index set of binary variables.
- $A \in \mathbb{R}^{m \times n}$ is the matrix of coefficients.
- $b \in \mathbb{R}^m$ is the vector of constraints.
- $c \in \mathbb{R}^n$ is the cost vector.

One common approach to solve a MILP is to relax the binary variables, allowing them to take values between 0 and 1 (Huang et al., 2025). This transforms the problem back to a linear form, which can then be solved using methods like the Simplex algorithm.

Relaxed Linear Programming Problem:

$$\begin{aligned} \text{Minimize} \quad & c^T x \\ \text{subject to} \quad & Ax \leq b \\ & 0 \leq x_i \leq 1, \quad \forall i \end{aligned}$$

If the relaxed binary variables take values of exactly 0 or 1 in the solution, the problem is solved optimally. However, if the binary variables end up with fractional values between 0 and 1, further steps are necessary to find the optimal solution. In this case, the branch-and-bound method may be applied. It will force the binary variable to take a value of either 0 or 1.

This process is likely how built-in solvers in Python handle MILP problems. While the exact algorithms used by some solvers are proprietary, the general approach typically involves branching to enforce the integer constraints on the binary variables and then relaxing the problem, solving it as a continuous linear problem.

2.5.2 Solving Linear Optimization Problems with the Simplex Method

To solve linear optimization problems using the Simplex method, the problem must first be reformulated into standard form. This involves converting all inequality constraints into equality constraints by introducing slack variables. These slack variables represent the unused portion of a constraint and are constrained to be nonnegative, just like the original decision variables.

For example, consider the inequality constraint:

$$x \leq b$$

By introducing a slack variable s, this inequality can be rewritten as an equality:

$$x + s = b$$

Using this approach, each inequality constraint is converted into an equality constraint by adding a slack variable. The problem setup should then have the following form:

$$\begin{aligned} \text{Maximize} \quad & w \\ \text{subject to} \quad & Ax = b, \\ & x \geq 0 \quad \forall x, \\ & x = (x_1, x_2, \ldots, x_n, s_1, s_2, \ldots, s_m) \end{aligned}$$

The feasible region defined by these constraints remains the same as in the original problem.

The Simplex algorithm exploits the fact that the optimal solution to a linear program, if it exists, is located at a vertex (corner point) of this region.

To begin, the variables are divided into two categories:

- **Basic variables**: Equal in number to the number of constraints.
- **Non-basic variables**: The remaining variables.

The number of non-basic variables corresponds to the degrees of freedom in the system, i.e., the number of variables that can be fixed while still having a solvable system of equations.

2.5.2.1 Initialization in the Simplex Method The Simplex method starts by selecting a basic feasible solution (BFS), which corresponds to a vertex of the feasible region. In simpler cases, this is achieved by setting all non-basic variables to zero, which provides a straightforward starting point if such a solution exists. However, this approach is not always feasible or optimal. More efficient initialization strategies are available. Examples of such methods are the Two-Phase method, Big-M method, and Warm Start technique.

Two-Phase Method and Big-M Method
Both of these methods involve adding artificial variables to the problem, making it easier to solve. The goal is to find a solution where the artificial variables are zero, making it a feasible solution to the original problem. Once a feasible solution is found, it is used as the starting point for the Simplex algorithm.

Two-Phase Method: The setup involves minimizing the sum of the artificial variables:

$$\text{Minimize} \quad (x, x_a) \quad 1m^T \cdot x_a$$

$$\text{subject to} \quad Ax + x_a = b$$

$$x_a, x \geq 0$$

A trivial solution, such as $x_a = b$ and $x = 0$, is a feasible starting point. The optimal solution is reached when the artificial variables x_a are zero. If no feasible solution is found, the original problem has no feasible solution.

Big-M Method: This method also uses artificial variables but with a slightly different formulation. The key difference is that the objective function includes a large penalty term for the artificial variables:

$$\text{Minimize} \quad (x, x_a) \quad c^T x + M \cdot 1m^T \cdot x_a$$

$$\text{subject to} \quad Ax + x_a = b$$

$$x_a, x \geq 0$$

where M is a large constant. This forces the algorithm to drive the artificial variables to zero. If a feasible solution is found with the artificial variables set to zero, it can be used as the starting point for the original problem.

Warm Start
The Warm Start technique is particularly useful when solving a series of similar optimization problems. It reuses the solution from a previous run as the starting point for a modified version of the problem. This can help place the algorithm closer to the optimal solution and reduce computational time (Huang et al., 2025).

With the setup for the initial point done, the method works as follows:

2.5.2.2 Steps of the Simplex Method:

1. **Initialization:** Initialize all non-basic variables the preferred method.
2. **Solve for basic variables:** Using the current basis, solve the system $Ax = b$.
3. **Check for improvement:** Identify a non-basic variable with the potential to increase the objective function value. This is done using the reduced cost (or the gradient of the objective function).
4. **Pivot:** Increase the selected non-basic variable until one of the basic variables becomes zero (i.e., hits a constraint).
5. **Update basis:** The entering variable becomes basic, and the leaving variable (the one that dropped to zero) becomes non-basic. Update the system accordingly.
6. **Repeat:** Continue iterating steps 3–5 until no further improvement can be made (all reduced costs are non-positive in a maximization problem).
7. **Termination:** At this point, where no further improvement can be made, the current basic feasible solution is optimal.

2.5.3 Branch and Bound for MILP Problems

The Branch and Bound (B&B) method is a widely used algorithm for solving Mixed-Integer Linear Programming (MILP) problems. The main idea is to systematically explore subsets of the feasible solution space by dividing the problem into smaller sub-problems (branches), forming a search tree.

The process begins by formulating the MILP problem and identifying the set of feasible solutions D. The original problem can then be relaxed by removing the integrality constraints, resulting in a standard Linear Programming (LP) problem. This relaxed problem can be solved using, for example, the Simplex method, which provides a lower bound on the objective value.

$$P = (D, f), \quad f : D \to \mathbb{R}$$

The optimal solution to the MILP is denoted x^*, representing the best feasible solution that minimizes the objective function f. The best solution found so far, $\hat{x}$, is stored as a global variable.

In each iteration, a sub-problem is selected from a queue of unexplored sub-problems $\mathcal{L}$. If a feasible solution x' within the sub-problem satisfies $f(x') < f(\hat{x})$, it replaces the current best. If no better solution exists within the sub-problem, it is pruned and removed from $\mathcal{L}$; otherwise, it is further branched into smaller sub-problems and added back to the queue. The algorithm terminates when $\mathcal{L}$ is empty, at which point the best solution $\hat{x}$ is returned (Huang et al., 2025).

The general structure of the B&B algorithm is presented above, but its efficiency relies on some key implementation choices. These include branching variable selection, node selection strategies, node pruning techniques, and the use of cutting planes to tighten relaxations. For an overview of these components, see (Huang et al., 2025).

2.5.4 Particle Swarm Optimization (PSO)

Particle Swarm Optimization (PSO) is a population-based heuristic optimization method introduced by Eberhart and Kennedy (1995). The method is inspired by the herd behaviour of birds and fish. It operates by iteratively improving a group of potential solutions, called particles, that are moving around in a space of potential solutions. It is often used to solve complex optimization problems where traditional methods, such as gradient search algorithms, are ineffective or hard to implement.

PSO is particularly effective in situations where the search space is large, the problem formulation is non-linear, or where an analytical description of the function gradient is unavailable. Common applications include machine learning, engineering design, economic modelling, and energy system optimization. Due to its flexibility, simple implementation, and ability to handle both continuous and discrete optimization problems, PSO has become a widely used tool across various fields (Poli et al., 2007).

The PSO algorithm has previously been used for optimization of both Energy Management Systems (EMS) and microgrid optimization. In this review by Thirunavukkarasu et al. (2022), the PSO algorithm is presented as one of the top five EMS optimization techniques. The PSO algorithm performs good or moderately on all of the criteria they listed.

Each particle in the PSO method has three attributes: position x_i, velocity v_i, and personal-best p_i. The position x_i describes what position the particle has in the solution space, i.e., its currently proposed solution. The velocity vector v_i describes the movement of the particle, and the personal best gives the best solution yet proposed by the particle. The PSO algorithm also keeps track of a global best, which is the best solution proposed yet through all iterations of all particles.

The first step of particle swarm optimization is to initialize the positions of the particles. Preferably, they are somewhat spread out in the potential solution space. They also need an initial velocity. The velocity for particle i is then iteratively updated as:

$$v_i(t+1) = wv_i(t) + c_1 r_1(p_i - x_i) + c_2 r_2(g - x_i)$$

where w is the inertia value that balances the exploration of the particle with the belief in global and self-best solutions. c_1 and c_2 determine how much the particle is drawn to the personal p_i, respectively, the global-best g. Lastly, r_1 and r_2 are two random numbers between 0 and 1.

The position x_i is also iteratively updated with the velocity as:

$$x_i(t+1) = x_i(t) + v_i(t+1)$$

An illustration of how the particle position is updated can be seen in Figure 3. After each iteration, the current solution is evaluated and compared to both the particle's self-best and the global-best. If the new solution is better, the self-best and/or global-best values are updated accordingly. This process continues until the algorithm converges, which happens either after a predetermined number of iterations or when the improvement of the algorithm is sufficiently small.

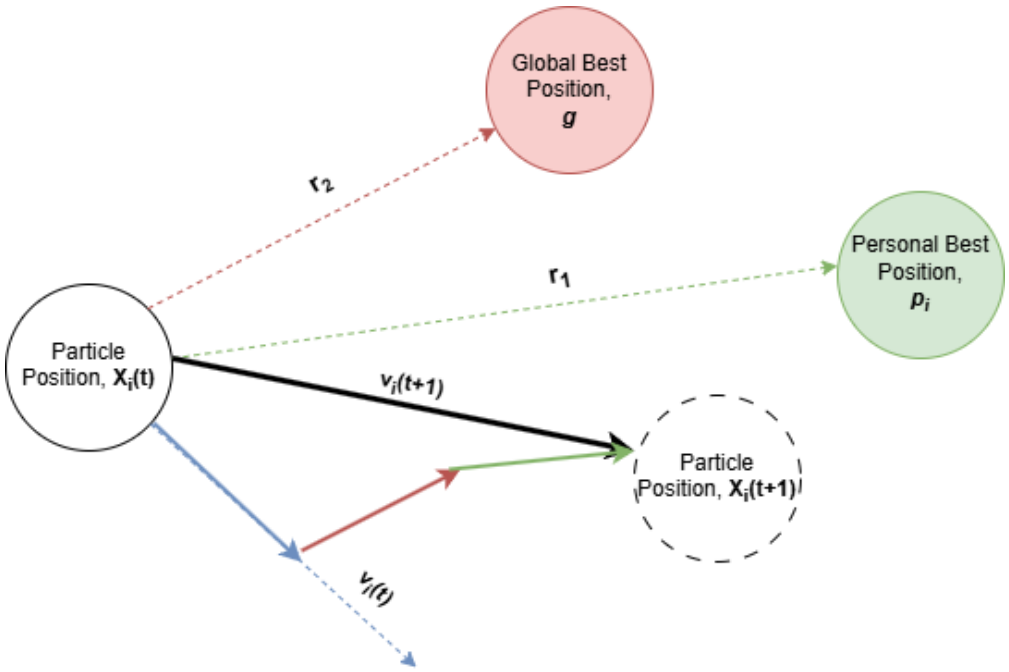

Figure 3: Visualization of the position update mechanism in the Particle Swarm Optimization (PSO) algorithm. Each particle adjusts its position in the solution space based on its current velocity, its personal best position, and the global best position found so far.

The optimal number of iterations is problem-specific. Too few iterations may lead to premature convergence, while too many give unnecessary computational complexity.

2.5.4.1 Inertia Weight, w Inertia weight, w, is one of the parameters in the PSO algorithm that affects the velocity of the particles. A large value of w ($w > 1$) indicates fast movement in the solution space, while a smaller value ($w < 1$) means slower movement and faster convergence of the particles. The inertia value is therefore a trade-off between exploring the whole solution space and achieving fast convergence.

The inertia value was initially introduced by Shi and Eberhart (1998) as a constant value. However, since the introduction of inertia values, many different approaches have been suggested to choose inertia weights. The article *Inertia Weight Strategies in Particle Swarm Optimization* by Bansal et al. (2011) reviews 15 different methods for choosing inertia weights and suggests two different methods that were the best suggestions for two different objectives.

The best method for high accuracy was the *Chaotic Inertia Weight*, and the best method for efficiency was the *Random Inertia Weight*. For this problem, efficiency will be more important. The random inertia weight will therefore be the one applied in this thesis, and it is chosen as:

$$w = 0.5 + \frac{\text{rand}}{2}, \quad \text{where rand} \sim U(0,1)$$

2.5.4.2 Cognitive and Social Weights, c_1 and c_2 The parameters c_1 and c_2 are the cognitive and social weights. They balance how much the particles rely on their own experiences versus how much they follow the global-best solution.

c_1 determines the extent to which a particle trusts its own past experiences. A large value of c_1 indicates high self-confidence. c_2 dictates how much a particle is influenced by the global-best. A large value of c_2 means the particle is more likely to follow the swarm.

Again, there is a trade-off between these values. A large c_1 and small c_2 promote exploration but may prevent convergence. Conversely, a large c_2 and small c_1 facilitate quick convergence but increase the risk of getting stuck in a local minimum.

The study *Empirical study of particle swarm optimization* by Shi and Eberhart (1999) suggests setting c_1 and c_2 equal, while their newer study *Comparing inertia weights and constriction factors in particle swarm optimization* (Eberhart and Shi, 2000) suggests modifying the velocity update formula with a constant K and choosing $c_1 + c_2 > 4$ to guarantee stable convergence. Another approach is to start with a high c_1 (promoting exploration) and gradually decrease it while increasing c_2 over iterations, encouraging convergence toward the end of the process. Generally, values of c_1 and c_2 between 1.5 and 2.5 are recommended, though their optimal values may vary depending on the specific problem.

2.5.4.3 Number of Particles Additionally, the number of particles must be determined. A higher number of particles increases the initial diversity of the swarm, provided that an effective initialization scheme is used, which can improve the algorithm's ability to explore the solution space. More particles can also make the algorithm reach a good result with fewer iterations. However, it also increases the computational complexity per iteration. Empirical studies suggest that a suitable number of particles typically ranges between 10 and 30 (Clerc, 2010), although the optimal number of particles is problem-dependent.

3 Method

The wind power production data and spot price data were provided by Bodecker Partners, while the solar irradiation data was obtained from SMHI (2025).

3.1 Simulation

Simulations are a useful tool for understanding possible future scenarios. Efficient microgrid operation one year ahead requires an estimation of how wind production, photovoltaic production, and spot prices may evolve over the following year. While it is nearly impossible to accurately forecast 8 760 hours into the future, simulations can provide a range of possible outcomes, helping to form a better understanding of future developments.

The purpose of these simulations is to generate hourly scenarios for an entire year ahead. While they do not serve as precise forecasts, they offer a representation of potential variations and trends. By analysing multiple simulated future paths, valuable insights into uncertainties can be obtained, enabling the optimized construction of the microgrid.

3.1.1 Photovoltaic Power Simulation

For the simulation of photovoltaic power production, the EMBB method was applied. Solar power production in SE2 data from 2015 to 2024 was available, however, the installed solar capacity in SE2 has increased significantly in recent years. This makes data from earlier years less reliable, as the method involves interpolating between the installed capacity at the beginning and end of each year. When there are substantial changes in capacity within a year, such interpolation may not accurately reflect the actual production. Another problem with solar power production data is that the available data only represents the production that is sent to the grid, meaning that the production of electricity that is used directly in the buildings which they are connected is not taken into account. To avoid these problems, solar irradiation data was used instead to be able to get more reliable data and enough samples to make a statistically sufficient EMBB simulation. The solar irradiation data was then recalculated to photovoltaic energy as;

$$E = \frac{G \cdot A \cdot \eta \cdot t}{10^6}$$

Where

- E = Produced energy (MWh)
- G = Irradiation (W/m^2)
- A = Photovoltaic panels area (m^2)
- η = Energy efficiency (%)
- t = Time period (h)

A simulation of a future year was created by choosing a day from the month that was to be simulated and then taking a block of eleven days and adding it to the series. Eleven days was chosen from the suggestion of $n^{\frac{1}{3}}$ in the section 2.4.1. Eleven days was also evaluated as a good choice since it kept the autocorrelation structure from the historical data for the simulated data. The simulated data has similar behaviour in mean and standard deviation to the historical data. The ACF and PACF of the simulated series were analysed and compared to historical data to ensure consistency in autocorrelation structure. The historical and simulated ACF and PACF can be seen in Figure 4. The simulated photovoltaic production behaved similarly to the historical data.

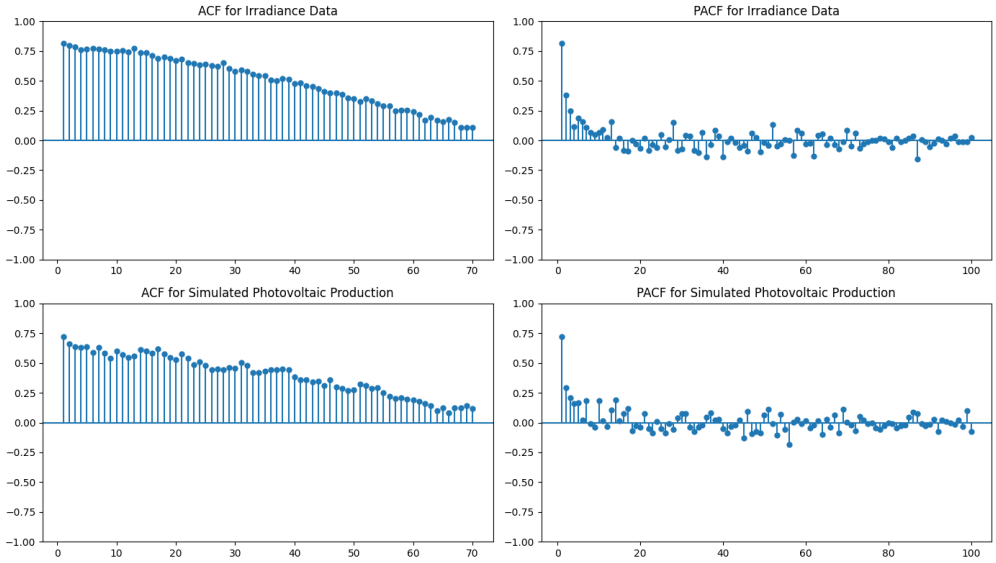

Figure 4: Comparison of the Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) between simulated photovoltaic power production and energy transformed irradiation data (2024).

The problem again with using this method is that no more extreme case than what has already happened can happen. Meaning that a really sunny winter or a sunfree summer can not happen, however, the sun has a rather stable cycle and the simulations were concluded as sufficiently good for this purpose. Two simulations, one for the lowest total production and one for the highest total production, can be seen in Figure 5. The total production is calculated as the sum of hourly generation over an entire year.

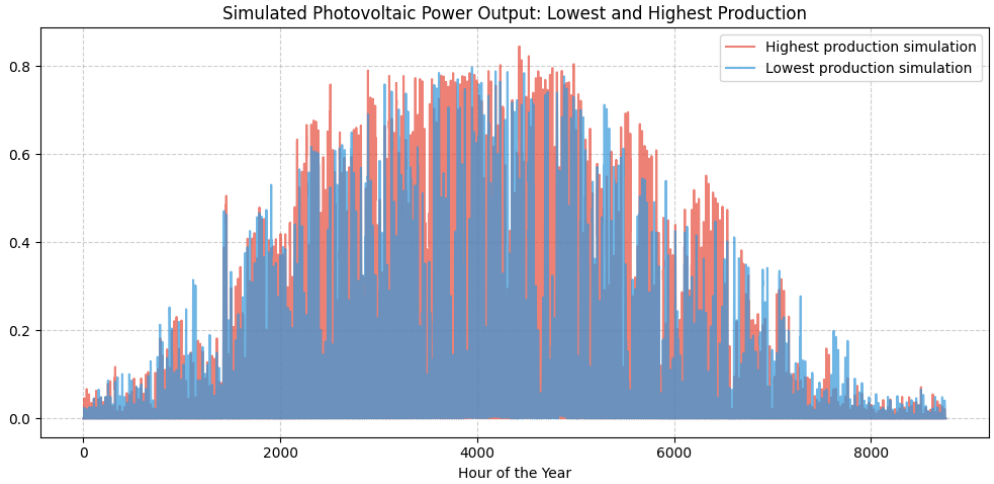

Figure 5: Simulated photovoltaic Power Production: Highest and Lowest Total Production Cases from 1000 Simulations.

EMBB was chosen due to the strong seasonal pattern in the solar irradiation data. Another method, such as stationary block bootstrap, suggested in Politis and Romano (1994) is also good for preserving autocorrelation structures, however, the seasonal pattern would be harder to capture. Other methods, such as SARIMA processes, didn't capture the behaviour as well. This is partly due to the problem with capturing the seasonality in sun hours per day, together with the daily and yearly seasons.

3.1.2 Wind Power Simulation

The same methodology used for the photovoltaic power simulation was applied to the wind power simulation. Wind production data from the period 2015–2024 was collected and normalized based on installed capacity. The normalized wind production for 2024 can be seen in Figure 6. Here, seasonal variations can be seen with higher production during the winter months. This can also be seen in Figure 15. To generate the simulated series, a random day was selected from the same month as the target simulation period, and a block of seven days was added to the simulated dataset.

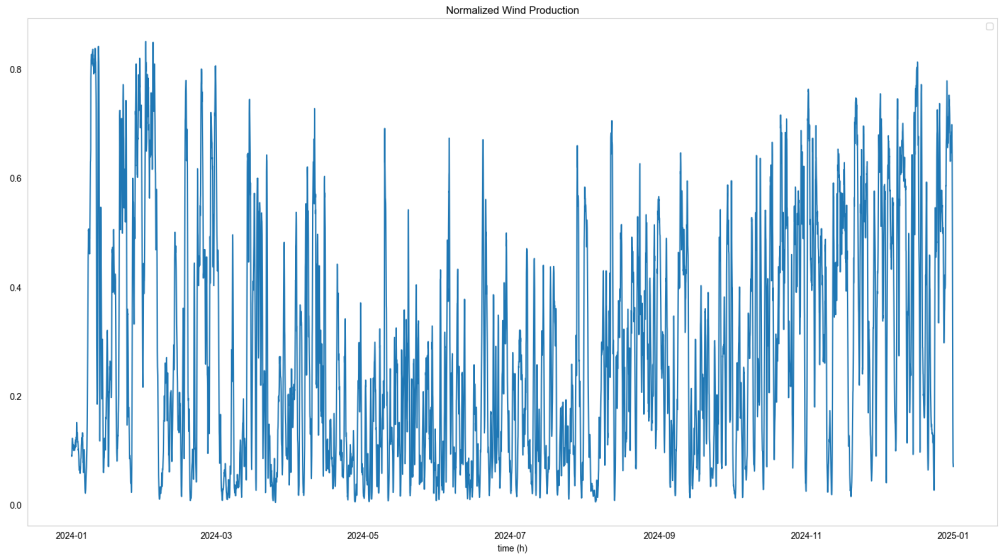

Figure 6: Wind production data in SE2 from 2024 that has been normalized with installed capacity.

This approach successfully maintained the mean and standard deviation of the original dataset. The choice of a seven-day block length was again chosen from the suggestion in Section 2.4.1. The choice of block length is a trade-off between introducing randomness and preserving the dependencies in the data. The preservation of the dependence structure can be seen in Figure 7.

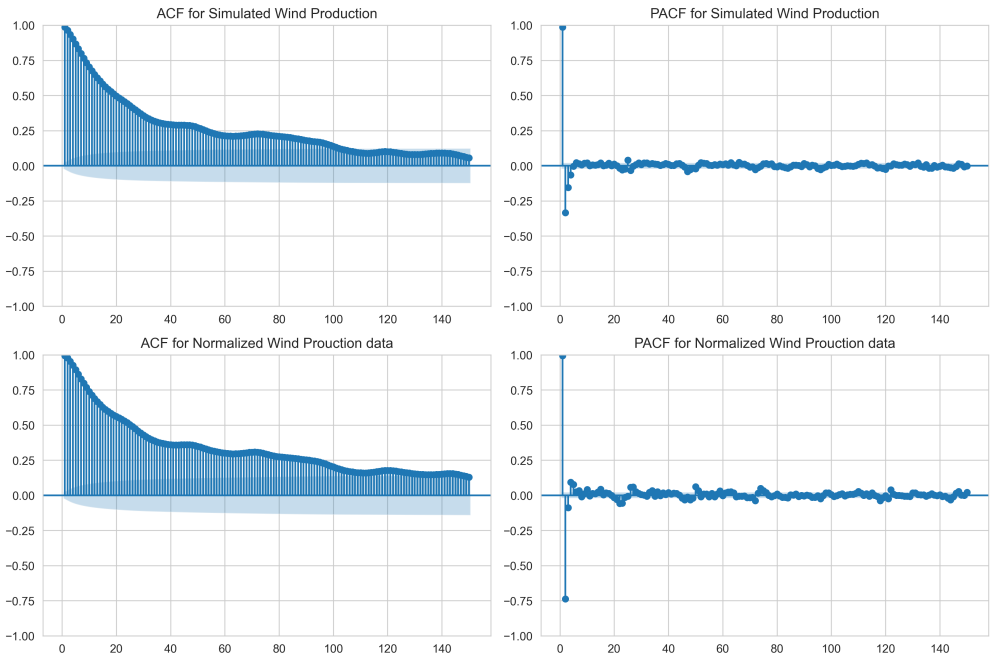

Figure 7: Comparison of the Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) between simulated Wind power production and normalized wind production data (2024/01/01-2024/12/31).

Similarly, as for the photovoltaic power simulation, a plot of the lowest total production simulation and for the highest total production simulation can be seen in Figure 8.

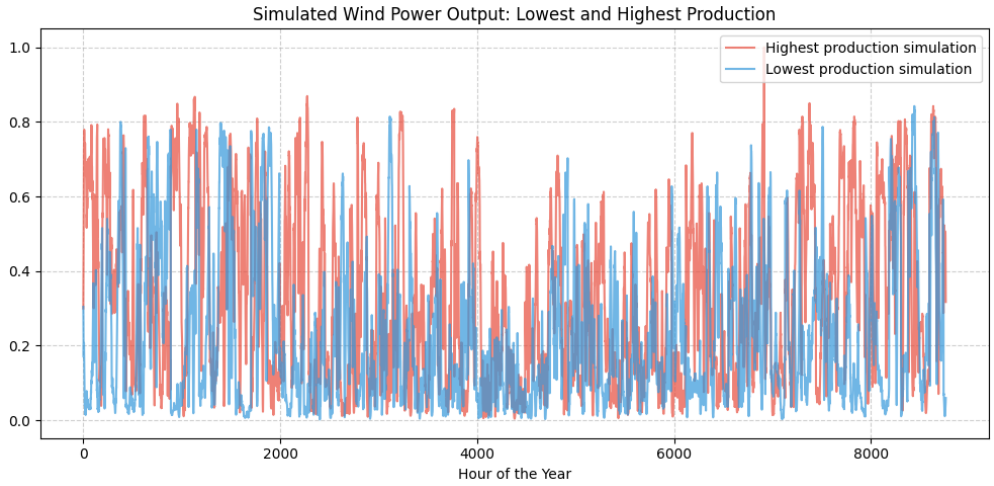

Figure 8: Simulated Wind Power Production: Highest and Lowest Total Production Cases from 1000 Simulations.

3.1.3 Construction of Spot Price Curve

Capturing the behaviour of the spot price proved particularly challenging. Various time series models, including ARX, SARIMA, and ARMA, were tested in an attempt to model the data. However, visual inspection, see Figure 9 suggested non-stationarity, and the models failed to produce reliable results. To further investigate the stationarity, an ARCH test was conducted, yielding a p-value numerically equal to zero. This result strongly indicates the presence of heteroscedasticity in the data, confirming that the variance is not constant over time. As a result, the spot price was not simulated directly. Instead, a representative spot price profile was constructed based on the simulated wind data.

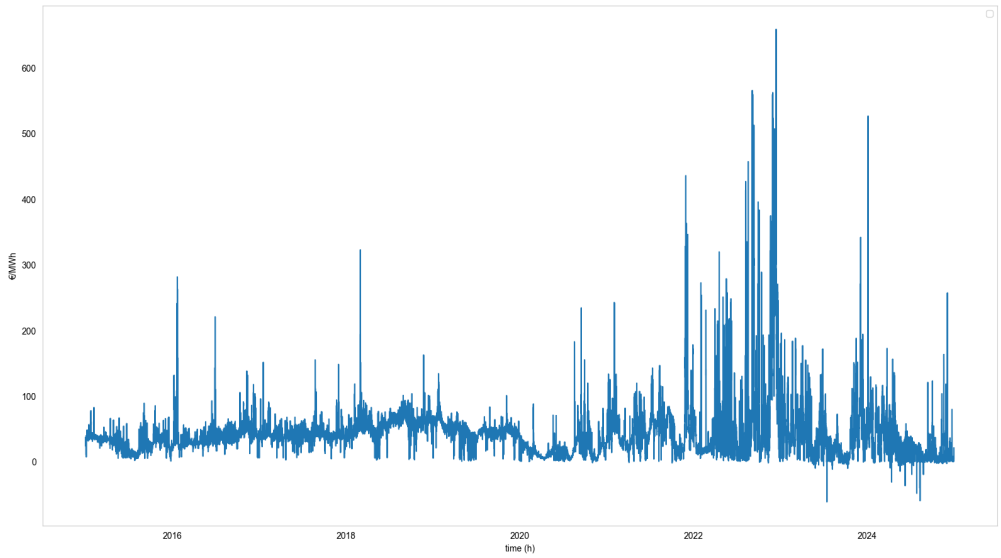

Figure 9: Hourly spot price data (2015-2024).

The primary objective of the simulations was to evaluate the sensitivity of the results under different renewable energy production scenarios. Therefore, the focus was placed on accurately simulating wind and photovoltaic production. Nevertheless, a realistic hourly spot price curve was also necessary to run the simulations properly.

The spot price simulation was inspired by the idea behind block sampling. That is, sampling consecutive chunks of data to preserve local structure and temporal dependencies. However, instead of sampling price blocks independently, the method is coupled with the wind production simulation. Specifically, the same days sampled in the wind simulation are used to extract the corresponding days from historical spot price data adjusted for inflation.

This approach not only retains the autocorrelation and structure of the spot price, as with traditional block sampling, but also captures the empirical correlation between wind production and spot price. For the location considered in this report (SE2), the mean correlation between spot prices and wind production during the period 2019-2024 was around -0.32 and a median of -0.39 (see Appendix, Table 14).

To validate the realism of the generated time series, autocorrelation (ACF) and partial autocorrelation (PACF) plots were examined. While the autocorrelation in the simulated data decays more rapidly compared to the real data, the seasons are still correct, and the seasonal patterns are reasonably well preserved. For the purposes of this study, the simulated spot price curve was therefore deemed sufficiently good, see Figure 10.

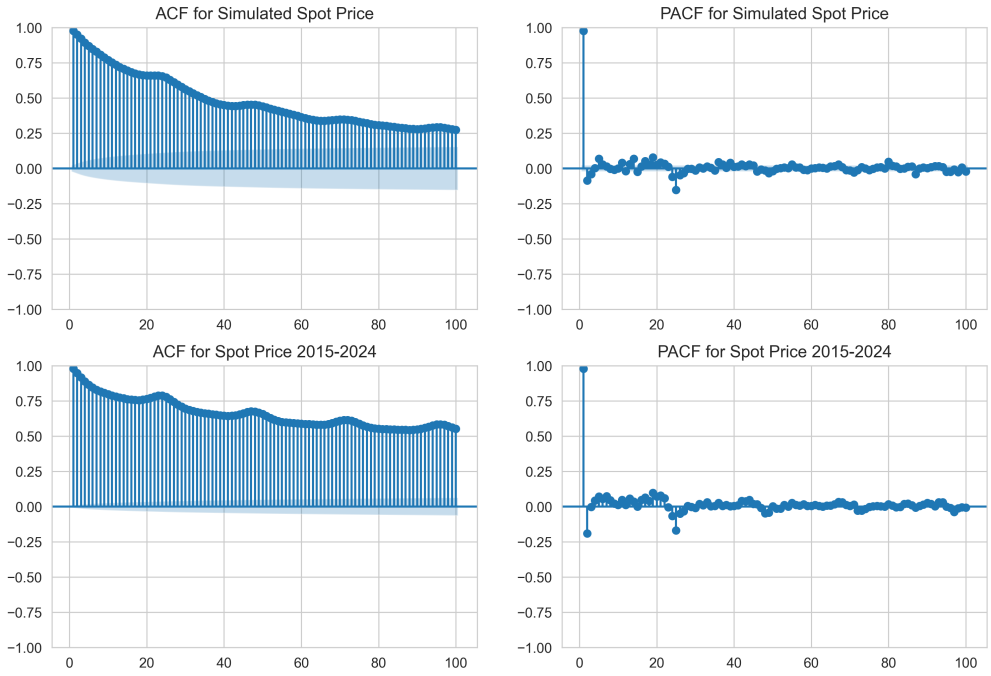

Figure 10: Comparison of the Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) between simulated spot price and spot price data (2015-2024).

Moreover, this methodology could easily be extended to include regulation and intra-day markets in future simulations or optimizations, as the same sampling logic can be applied to those price series as well.

3.2 Optimization

The aim of the optimization was to construct the optimal configuration of photovoltaic power, wind power, and batteries and to run the microgrid in the most efficient way. This was to be done for one year ahead. The problem was made a bit simpler by applying perfect foresight for the following day, meaning that all spot prices and production were known for all 24 hours the next day when the bids to the day-ahead market were placed.

3.2.1 Parameter Definitions & Set Up

In the following section, the parameter values for wind turbines, batteries, and diesel generators were provided by Bodecker Partners. The values for photovoltaic panels are sourced from HemSol (2025).

Before the optimization could begin, it was necessary to define several key parameters such as lifespan, capacity, investment costs, and operational costs for both wind turbines and photovoltaic panels. These values depend on the specific models used and can easily be changed if needed. The parameter values in this project are shown in Table 1.

	Wind Turbines	PV Panels
Life span (years)	35	25
Maximum capacity (MWh)	7	1
Investment cost (EUR)	10 000 000	13 636
Maintenance cost (EUR/year)	147 000	1 364

Table 1: Parameters for **Wind Turbines and Photovoltaic Panels**.

Lithium-ion batteries were chosen for energy storage. These were modelled with a maximum capacity of 1 MWh and assumed to be able to fully charge or discharge in one hour. They could operate between 20% and 80% of their full capacity and had a lifetime limit of 15 000 cycles. A charging efficiency of 87.5% was used to account for energy losses, see Table 2.

	Lithium-Ion Battery
Maximum capacity (MWh)	1
Investment cost (EUR)	450 000
Operational cost (EUR/year)	0
Life span (year)	20

Table 2: Parameters for **Lithium-Ion Battery**.

If the microgrid and market could not meet the baseload, a diesel generator would serve as backup. Since it is assumed to already exist at the data center, it was excluded from the investment analysis. Only the operational cost of diesel fuel was considered. The associated cost is presented in Table 3.

	EUR/MWh
Fuel cost	317

Table 3: Operational cost of **Diesel generator**.

To obtain results, assumptions about baseload and grid limitations were required. These values can easily be adjusted, but the specific ones used to produce the results below are presented in Table 4.

	Value (MW)
baseload	20
Grid import limitation	18

Table 4: Assumptions for **Baseload and Grid Limitation**.

The data used to optimize the microgrid configuration included spot prices, wind production, and solar irradiation converted to photovoltaic power generation from 2024 in SE2. Perfect foresight was assumed during operation, meaning that spot prices as well as wind and photovoltaic production for the following day were known in advance. Additionally, the costs of the land for the microgrid were ignored, and it was also assumed that the microgrid would have no effect on overall market prices.

To find the optimal configuration, the grid export limitation was initially set to zero. This serves two main purposes. The first is to prevent the optimization from producing unrealistic results. Since the setup assumes perfect foresight one day in advance, the system could otherwise exploit energy trading to generate infinite profits, which is not realistic. Disabling energy sales to the grid helps avoid this issue. The second purpose is to focus on the primary objective of the microgrid: to cover the baseload. By not allowing energy to be sold, the optimization identifies the minimum system size required to meet the baseload demand.

However, to provide a more accurate representation of the operational costs, a second optimization was carried out for the identified optimal configuration, this time allowing energy for 200 MW to be sold to the grid. This gives a fairer comparison of the different cost components associated with operating the microgrid.

3.2.2 Optimization Algorithm

The optimization process is divided into two main components: the PSO algorithm and the fitness function, which includes a nested forecast function. These components work together to evaluate and identify the most cost-effective microgrid configuration.

The fitness function simulates the operation of the microgrid over an entire year and calculates the associated total cost. This cost includes both fixed investment costs and operational costs, based on how the microgrid is utilized throughout the year. The operational strategy is determined by the forecast function, which looks at production and spot prices to optimize system usage.

In the PSO algorithm, each particle represents a unique microgrid configuration. Each particle is evaluated using the fitness function, and its corresponding cost serves as the particle's fitness score. The PSO algorithm then iteratively searches for the configuration that minimizes this cost.

3.2.2.1 Fitness Function The fitness function is the objective function to be minimized, evaluating the cost-effectiveness of a given microgrid configuration over a year. It consists of an **outer layer**, which tracks the system's operational and investment costs, and an **inner layer**, which determines the optimal operation of the microgrid for the following day using a function called the forecast function. This includes scheduling energy production, battery usage, and market interactions. This forecast function runs daily at 11:00 AM, aligning with the deadline for placing spot orders on the Nord Pool market.

The fitness function takes as input a proposed configuration consisting of the number of wind turbines, photovoltaic panels, and batteries. At the end of each simulation, it calculates the total yearly cost, which includes investment, maintenance, and operational costs. This total cost serves as the objective function value that the optimization algorithm seeks to minimize.

A flowchart illustrating the function's process can be seen in Figure 11.

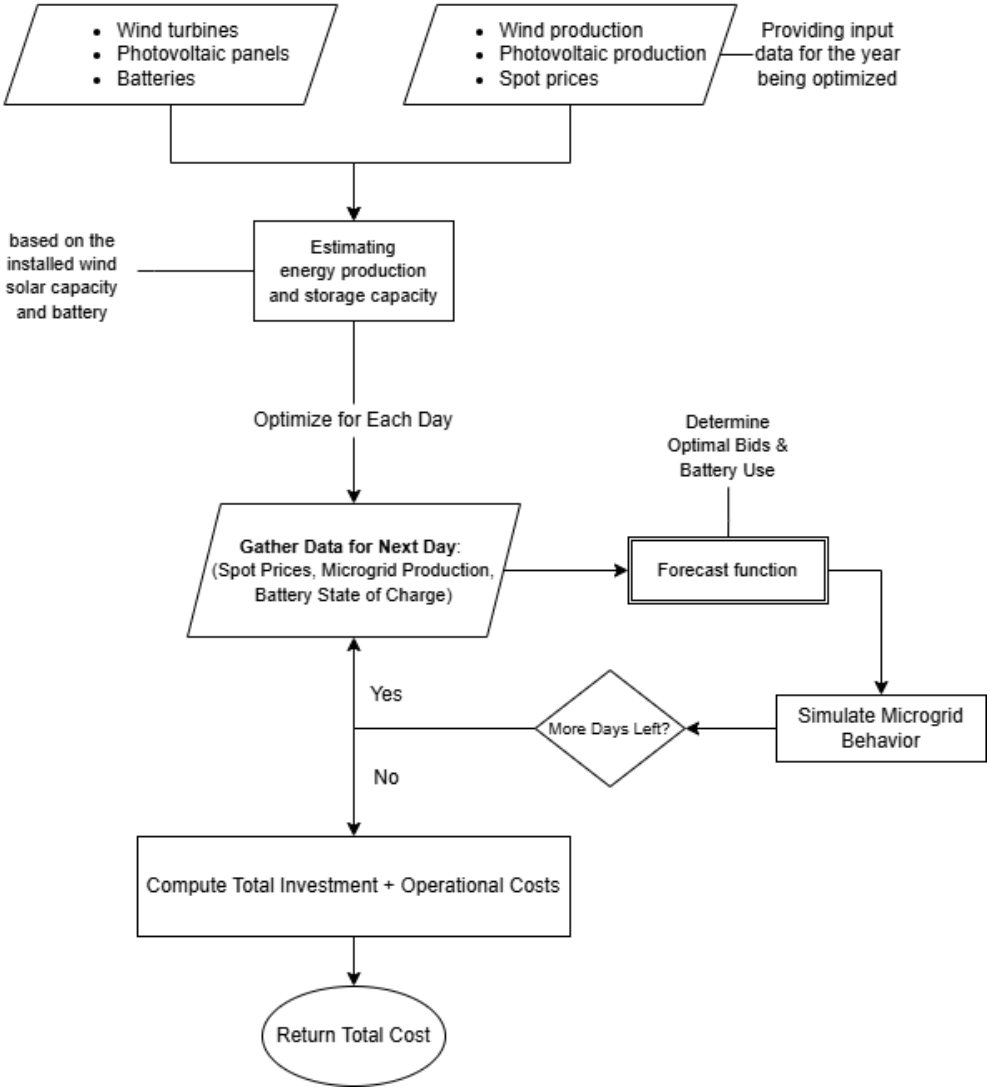

Figure 11: Flowchart illustrating the fitness function

3.2.2.2 Forecast Function Within the fitness function, an iterative loop calls a 'forecast function' for each day to determine the optimal operation of the microgrid for the following day. This includes deciding which spot market orders to place, how the battery should be charged or discharged, and whether any production should be curtailed. The function takes as input the microgrid's expected production and spot prices for the next day, which are known in advance due to the assumption of perfect foresight. Based on these inputs and a set of constraints, the function optimizes the system's actions for the upcoming day.

The optimization is implemented in Python using the CVXPY library, which provides a framework for defining and solving convex optimization problems. The objective function is minimized using the SCIP solver from Achterberg (2009), which integrates a combination of the simplex algorithm and branch-and-bound techniques to efficiently handle both continuous and discrete variables. SCIP is integrated directly into CVXPY and does not require any external installation.

The optimization problem is formulated for managing a microgrid on an hourly basis for the upcoming day. All input data is assumed to be known in advance under the assumption of *perfect foresight*.

Input Parameters (given for each hour of the next day):

- Spot market price
- Total production (wind and photovoltaic)
- Baseload demand
- Grid import limit (maximum allowed purchase from the market)
- Grid export limit (maximum allowed sale to the market)
- η_{ch} - charging efficiency
- η_{dis} - discharging efficiency

Decision Variables:

- **Charge (C)** – energy charged to the battery [MWh]
- **Discharge (D)** – energy discharged from the battery [MWh]
- **Buy** – electricity bought from the spot market [MWh]
- **Sell** – electricity sold to the spot market [MWh]
- **State of Charge (SoC)** – battery energy level at each hour [MWh]
- **Slack** – soft constraint variable allowing deviation from baseload demand
- **Curtailment** – share of renewable production curtailed (between 0 and 1)
- $Z_{\textbf{battery}}$ – binary variable ensuring charge/discharge exclusivity

Constraints
In the following constraints, t refers to time and T is the set of the 24 hours of the following day.

Battery-related constraints:

- **State of Charge (SoC) dynamics:** The battery's SoC at time t depends on the previous time step and the net energy charged or discharged, accounting for efficiency:

$$\mathrm{SoC}_t = \mathrm{SoC}_{t-1} + \eta_{\mathrm{ch}} \cdot C_t - \frac{D_t}{\eta_{\mathrm{dis}}}, \quad \forall t \in T,\, t > 0$$

- **SoC limits:** The state of charge must always lie within its physical bounds:

$$\mathrm{SoC}^{\min} \leq \mathrm{SoC}_t \leq \mathrm{SoC}^{\max}, \quad \forall t \in T$$

- **Charging limitation:** The battery can only be charged with energy that is physically available and within storage capacity:

$$\begin{aligned} &C_t \leq \mathrm{Production}_t + \mathrm{Buy}_t \\ &\eta_{\mathrm{ch}} \cdot C_t + \mathrm{SoC}_{t-1} \leq \mathrm{SoC}^{\max} \quad \forall t \in T \end{aligned}$$

- **Discharging limitation:** The battery cannot discharge more energy than is available above the minimum SoC threshold:

$$\frac{D_t}{\eta_{\mathrm{dis}}} \leq \mathrm{SoC}_{t-1} - \mathrm{SoC}^{\min}, \quad \forall t \in T$$

Microgrid-related constraints:

- **Baseload prioritization:** The total energy available must cover the baseload demand, allowing a slack variable for feasibility:

$$\mathrm{Production}_t - C_t + D_t + \mathrm{Buy}_t - \mathrm{Sell}_t + \mathrm{Slack}_t \geq \mathrm{Baseload}_t, \quad \forall t \in T$$

- **Net energy balance:** Net energy is defined as the total incoming energy minus the total outgoing energy at each time step:

$$\mathrm{NetEnergy}_t = \mathrm{Production}_t + \mathrm{Buy}_t + D_t - (\mathrm{Sell}_t + \mathrm{Baseload}_t + C_t)$$

  To avoid an energy surplus in the system, net energy must not be positive:

$$\mathrm{NetEnergy}_t \leq 0$$

  Any excess production should be either stored, sold, or curtailed to maintain this balance.

- **Curtailment constraint:** Curtailment must be a value between 0 and 1 (0%–100%):

$$0 \leq \text{Curtail}_t \leq 1, \quad \forall t \in T$$

- **Grid import constraint:** The net grid import must not exceed the allowed import capacity:

$$\text{Buy}_t \leq \text{GridLimitBuy}_t, \quad \forall t \in T$$

- **Grid export constraint:** The net grid export must not exceed the allowed export capacity, considering that excess energy is sold:

$$\text{Sell}_t \leq \text{GridLimitSell}_t \qquad \forall t \in T$$

- **Charge/Discharge Exclusivity:** To prevent simultaneous charging and discharging, a binary control variable z_t is used:

$$\begin{aligned} &\eta_{\text{ch}} \cdot C_t \leq z_t \cdot \text{SoC}^{\max} \\ &\frac{D_t}{\eta_{\text{dis}}} \leq (1 - z_t) \cdot \text{SoC}^{\max} \end{aligned} \quad \forall t \in T$$

Objective Function

$$\max \sum_{t \in T} \Big(\text{Sell}_t \cdot \text{Price}_t - \text{Buy}_t \cdot \text{Price}_t - M \cdot \text{Slack}_t + N \cdot (1 - \text{Curtail}_t) \Big)$$

where M and N are arbitrary large numbers and $M \gg N$

From this problem setup, the optimization is carried out using the SCIP solver. This allows the model to navigate the feasible space and identify an optimal daily operation strategy for the microgrid. By aligning economic incentives with operational constraints, the optimization ensures cost-effectiveness, reliability, and sustainable energy management on a day-ahead basis.

3.2.2.3 Applying the PSO Algorithm to Generate a Solution The PSO algorithm seeks the best possible combination of system components (wind turbines, photovoltaic panels, and batteries) by adjusting the particles based on their fitness values. Through iteration and refinement over multiple generations, the optimal configuration that minimizes the total annual costs while meeting the microgrid's energy requirements is found.

In the theory section, different parameters and choices of their values are discussed for the PSO algorithm. The optimization is done over a year on an hourly basis. This makes the efficiency of the code highly important. The inertia weight that was best for high efficiency was therefore chosen, and it was

$$w = 0.5 + \frac{\text{rand}}{2}, \quad \text{where rand} \sim U(0,1)$$

The social c_1 and cognitive c_2 weights were chosen to balance the exploration with a stable convergence. Therefore, the cognitive weight was chosen as 2.5 to begin with, and then it slowly decreased to 1.5. The social weight, on the other hand, starts with a lower value of 1.5 and then slowly increases to 2.5. This way of choosing the social and cognitive weight promotes exploration in the beginning while it then slowly lets the particles be more strongly affected by the global optimal, which improves the convergence. The convergence is further improved by the choice of c_{min} and c_{max}, as they ensure that $c_1 + c_2$ never falls below 4. The formulas used to calculate c_1 and c_2 can be seen below.

$$c_1(t) = (c_{1,\min} - c_{1,\max}) \cdot \frac{t}{n_t} + c_{1,\max}$$
$$c_2(t) = (c_{2,\max} - c_{2,\min}) \cdot \frac{t}{n_t} + c_{2,\min}$$

where $c_{1,\max} = c_{2,\max} = 2.5$ and $c_{1,\min} = c_{2,\min} = 1.5$. t is the current time step, and n_t is the total number of iterations.

Additionally, both the number of particles and the number of iterations need to be selected. As mentioned in the theory section, a commonly recommended range for the number of particles is between 10 and 30. This served as a starting point for tuning the algorithm to the specific problem at hand. After testing different configurations, the final number of particles was set to 40.

A similar approach was used to determine an appropriate number of iterations. Through empirical testing, 40 iterations were found to provide a good balance between computational cost and performance, avoiding premature convergence while still keeping the runtime reasonable.

4 Results

The result for the setup presented in Section 3.2.1 was calculated for three different requirements on baseload coverage: 80, 90, and 100%, with 80% meaning that it is guaranteed that the baseload will be covered 80% of the hours of the year. For the resulting number of wind turbines, photovoltaic panels, batteries, and total installed capacity, see Table 5.

Baseload Guarantee	Wind Turbines	PV Panels	Batteries	Total Capacity (MW)
80%	3	3	0	24
90%	4	4	1	33
100%	17	3	11	133

Table 5: Installed capacity in different microgrid configurations under varying baseload guarantees.

The percentage given in the Table above is the minimum percent of the hours required to cover. However, the actual percentage covered differed from the minimum requirement in one of the three cases. The actual baseload coverage is presented in Table 6.

Baseload Guarantee	Baseload Coverage
80%	82%
90%	90%
100%	100%

Table 6: Comparison of baseload coverage in microgrid configurations.

Table 7 presents the amount of energy sold and bought from the grid in total over the year. It also shows the average purchasing price and average sale price.

Baseload Guarantee	Energy Purchased (MWh)	Avg. Purchase Price (€/MWh)	Energy Sold (MWh)	Avg. Sale Price (€/MWh)
80%	120 168	28.3	0	-
90%	103 047	30.3	455	61.2
100%	34 258	40.5	152 898	12.6

Table 7: Energy purchased and sold on the spot market under different baseload guarantee levels.

Lastly, Table 8 shows the breakdown of the costs. The investment, the fixed yearly costs, and the operational costs. Fixed costs represent the costs of the investments divided by their expected life span plus the yearly maintenance cost. The operational costs are the spot orders and the usage of the diesel generator. The levelized cost of electricity (LCOE) is presented and calculated as the total annual cost divided by the annual energy consumption.

Baseload Guarantee	Operational Cost (M€/year)	Fixed Cost (M€/year)	Total Investment (M€)	LCOE (€/MWh)
80%	3.7	1.4	34	29
90%	3.3	1.9	45	30
100%	-0.54	7.7	178	41

Table 8: Cost comparisons between microgrid configurations.

5 Discussion

In Table 5, the microgrid configurations for the three baseload guarantee cases are presented. It can be observed that the result is relatively stable between 80% and 90% baseload coverage guarantee. The difference between the two cases is that the 90% case includes one extra unit of each microgrid component, increasing the installed capacity by 9 MWh and thereby enhancing total energy production. The additional battery also improves the reliability of the energy supply by helping to meet the baseload during periods of low renewable generation. In Table 6, it is shown that the actual achieved baseload coverage is 82% for the 80% case, while the 90% and 100% scenarios exactly meet their respective baseload requirements.

Raising the baseload requirement to 100% significantly alters the optimal configuration. The number of wind turbines increases considerably. Although wind turbines require a higher initial investment cost compared to photovoltaic panels, they provide a more consistent energy output throughout the year. While wind production is still seasonal, it continues both during winter and at night, unlike photovoltaic panels. Therefore, in order to cover 100% of the baseload at all times, the number of wind turbines increases, while the number of photovoltaic panels remains consistent with the 80 and 90% cases. The number of batteries also increases significantly from 1 to 11 units. Batteries are the one energy source that can provide the data center with energy during hours with low or no renewable generation. This explains their necessity in the 100% baseload case to guarantee uninterrupted energy delivery at all times.

Table 7 presents the microgrid's trading behaviour on the spot market under different baseload guarantee scenarios. For the 80% and 90% cases, the microgrid sells almost no excess energy to the market. Their installed capacities of 24 MW, respectively 33 MW, are barely sufficient to meet the 20 MW baseload during periods of optimal renewable generation. In contrast, the 100% case has a significantly higher installed capacity of 133 MW, enabling the microgrid to meet the baseload in many more hours and thereby substantially increase the volume of energy sold on the spot market. The amount of electricity purchased from the spot market remains relatively stable in the 80% and 90% cases, while it drops considerably in the 100% case. This aligns with the substantial increase in installed capacity, which reduces dependence on external supply.

As expected, the 100% case shows a clear spread between the average purchasing and selling prices. This is driven by the dynamics of the demand curve, as discussed in Section 2.2. When renewable generation is high, i.e., when the microgrid is producing excess electricity, other renewable energy producers are likely generating at high levels as well, which depresses the market price. Consequently, the selling price is typically low. Conversely, during low production periods, reduced renewable output across the grid increases reliance on more expensive energy sources, pushing market prices higher, resulting in higher purchasing prices.

Historical data support this pattern. Between 2019 and 2024, the average correlation between wind production and the spot price was approximately -0.32, see Table 14. As expected, since wind production is larger during winter and photovoltaic production is larger during

summer in SE2, wind and photovoltaic production are negatively correlated with an average of -0.24 between 2019 and 2022, see Table 16. Meanwhile, the correlation between solar production and the spot price has been modestly positive, around 0.028 during the period 2015–2022, see Table 15. The high number of wind turbines in the 100% scenario thus strengthens the negative correlation between the microgrid's total production and the spot price, which amplifies the spread between purchasing and selling prices.

Unexpectedly, the 90% case exhibits an inverted price spread, where the average selling price exceeds the average purchasing price. Although the microgrid in this scenario rarely sells energy to the spot market, it appears to take advantage of favorable market conditions when it does. This outcome may be attributed to the assumption of perfect foresight in the model, but also to efficient use of the battery system. By storing low-cost renewable energy during periods of high generation and strategically discharging it during hours of low production or high market demand, the microgrid can achieve a higher selling price than the average purchase price.

In Table 8 it can be seen that the fixed costs and investment costs are significantly increased in the 100% case. This is expected since the installed capacity increased by several hundred percent compared to the 80% and 90% case. The installed power capacity is also relatively higher in wind compared to photovoltaic in the 100% case. The wind turbines are more expensive compared to the photovoltaic panels, increasing both the fixed costs and the investment costs. Interestingly, the operational cost turns out to be negative, indicating that the microgrid earns money by selling energy to the grid. However, these relatively small profits are overshadowed by the much higher fixed and investment costs. This is also reflected in the LCOE, with the 80% and 90% cases having similar costs, while the 100% case is approximately 10 euros more per MWh.

For comparison, from an economic perspective, the microgrids could be compared to the alternative of operating without one. For this, assuming that the data center's entire baseload was to be covered by buying electricity from the spot market (with a full 20 MW grid connection). The annual operational cost would be approximately 4.3 million euros. In comparison, the scenarios with 80%, 90%, and 100% baseload coverage using a microgrid increase the annual cost by approximately 0.8 million euros, 0.9 million euros, and 2.9 million euros, respectively.

It should be noted that the forecast function assume perfect foresight in trading, which likely overestimates the real-world profitability of energy trading strategies. In reality, forecasting errors and increased complexity of having to deal with the intra-day and regulating market could significantly impact actual outcomes.

In addition to the numerical results, some Figures offer further insight. Figure 12 shows the diesel generator output for the scenarios: 80% and 90%. It also highlights the hours during which the baseload could not be covered by microgrid production or spot market purchases. The figure illustrates that the microgrid fails to meet the baseload during various hours throughout the year. A visual inspection reveals that there are no specific seasons during which the diesel generator is completely unused.

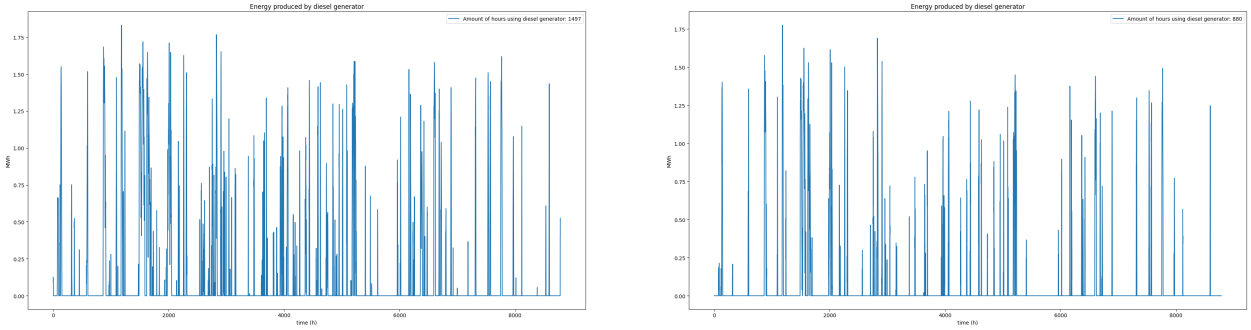

(a) 80% case

(b) 90% case

Figure 12: Yearly diesel generator production for microgrid configurations with 80% and 90% baseload guarantee.

Figure 13 presents the spot market transactions for the three scenarios. Positive values represent purchases from the grid, which are capped at 18 MW, the maximum amount that can be bought. Negative values represent sales to the grid. The selling limit is 200 MW, which is never reached, indicating that production is never forced to curtail.

The spot market transaction plot for the 80% and 90% scenarios shows increased electricity purchases during the summer months and reduced purchases during the winter. This can be explained by the relatively larger share of installed wind capacity, which typically produces more electricity during winter due to seasonal wind patterns. As a result, less electricity needs to be bought in the winter. In contrast, during summer, wind production is lower, leading to a higher need for purchases on the spot market. This increased demand is partly, but not fully, offset by photovoltaic production. In the 100% scenario, there is a significant increase in installed wind power capacity. The seasonal pattern of higher wind production during the winter is reflected in the spot market transactions, with more electricity being sold during the winter months compared to the summer. This is due to the surplus energy generated when wind production is high.

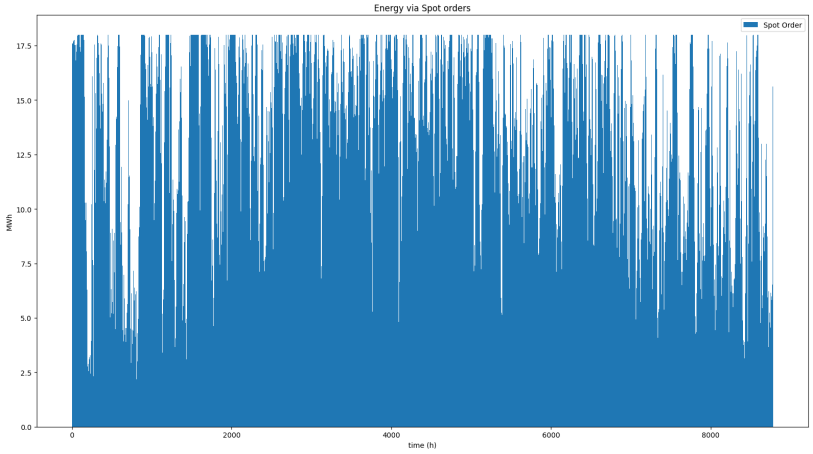

(a) 80% case

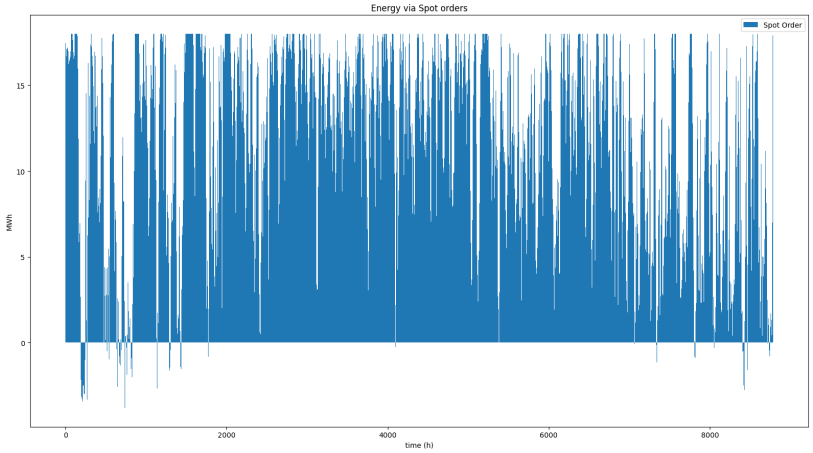

(b) 90% case

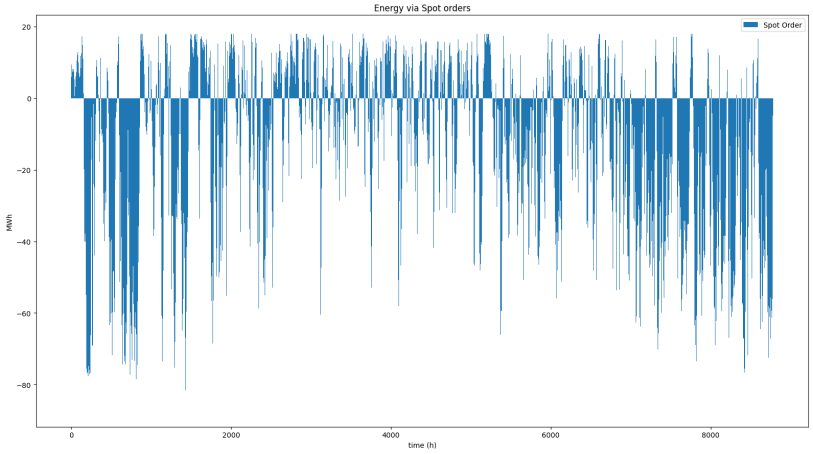

(c) 100% case

Figure 13: Yearly spot market orders for microgrid configuration in the 80%, 90%, and 100% baseload guarantee cases. The grid import limit is set to 18 MW and grid export limit is set to 200 MW

The microgrid may appear relatively large, with the smallest installed capacity reaching 24 MW. Considering that the system only needs to supply 2 MW to meet the baseload, this might seem oversized. However, this is likely due to the relatively low actual energy output from wind and photovoltaic power. Based on data from the optimized year, the average efficiency is around 23% for wind turbines and 9% for photovoltaic panels.

Another reason for the large microgrid sizes is likely the seasonal asymmetry in both wind and photovoltaic production. To consistently meet the baseload requirement throughout the entire year, including periods of low generation, a higher installed capacity is needed to ensure sufficient output across all seasons. The average monthly production for wind and photovoltaic power is shown in Figure 15 and Figure 16, respectively. This helps explain why the wind park tends to grow larger than the photovoltaic park when a strict 100% baseload coverage constraint is applied. Wind production is significantly more stable throughout the year, making it a more reliable source for consistently meeting baseload demand, despite its higher investment cost compared to photovoltaic panels.

Discussion of the Optimization

Since the optimization is performed on an hourly basis over an entire year, the process is computationally demanding and time-consuming. To improve efficiency, several acceleration techniques were explored. For example, attempts were made to set initial values for the Simplex method based on some of the initialization strategies discussed in Section 2.5.2.1. However, these efforts did not yield significant improvements in runtime, likely because the built-in solver already employs effective methods for initialization.

Another approach to reduce computational load was the application of stratified sampling, where samples are drawn from predefined subgroups, in this case, periods. Practically, this was implemented by randomly selecting some representative days from each month and assembling them into a synthetic, shortened year. This condensed dataset could then be used for preliminary optimization, helping to identify a promising region of the solution space in which the true optimum is likely to be found.

The restriction to SE2 applies only to the collected data. SE2 is the bidding zone in Sweden with the highest wind power production and is also one of the most affordable zones in terms of electricity prices, making it an attractive location for a data center. However, the same analysis conducted in this project can easily be replicated for another bidding zone simply by substituting the historical SE2 data with data from the desired zone. Similarly, all assumed parameter values can be adjusted to match specific case requirements.

5.1 Scenario Analysis

5.1.1 Full Baseload Coverage from the Grid

An interesting scenario is when the grid import limit is equal to the data center's baseload requirement. In this case, the data center does not require additional energy generation from, for example, a microgrid. However, it may still be economically beneficial to invest in a microgrid. By allowing the grid connection to match the data center's baseload demand while keeping all other parameters unchanged, the optimization resulted in no investment in a microgrid, meaning no wind turbines, PV panels, or batteries were installed. This indicates that, under current assumptions, purchasing the entire baseload demand from the spot market is more cost-effective than building a microgrid.

Nonetheless, this result is sensitive to several cost parameters. Modifying any of these, such as the investment cost of renewables or energy storage, could lead to a different outcome. As renewable technologies continue to evolve and become more cost-efficient, the economic attractiveness of local microgrid solutions is likely to increase. The same applies to battery systems; as discussed in the theory section, the battery market is rapidly expanding, and ongoing advancements may significantly improve the profitability of battery usage.

5.1.2 Lower Grid Connection Scenario

The national electrical grid is experiencing increasing pressure, and in the near future, this may lead to significant challenges for data centers. In such scenarios, data centers might not be prioritized over other critical societal services. Therefore, a scenario was simulated in which the data center (with 20 MW baseload) can only receive 10 MW from the national grid, equivalent to 50% of its baseload demand. While this is a more extreme case, it is not an unlikely one. The resulting system configurations are shown in Table 9. In all three cases, the actual baseload coverage matched the baseload coverage requirement exactly.

Baseload Guarantee	Wind Turbines	PV Panels	Batteries	Total Capacity (MW)
80%	11	21	10	108
90%	14	32	38	168
100%	100	14	33	747

Table 9: Microgrid configurations under varying baseload guarantees with 20MW baseload and 10 MW grid connection.

Even in the lowest baseload coverage requirement (80%), the microgrid has a total capacity more than 10 times the energy that needs to be covered by the microgrid. This is in line with the original result (see Table 5), where the energy to be covered by the microgrid was 2 MW, which also resulted in approximately 10 times more installed capacity in the 80% baseload coverage case, similarly for the 90% coverage cases. For the 100% case, the solution once again becomes extremely large compared to the lower baseload requirements, and the total capacity becomes unreasonably high.

Table 10 presents the total amount of energy bought from and sold to the grid over the course of a year in this scenario. It also reports the average purchasing and selling prices. As expected, there is again a price spread between the purchasing and selling prices observed across all configurations. The average purchasing price in the 100% baseload case is notably high. This is likely due to the large installed capacity, which means the data center only needs to purchase energy during hours of extremely low renewable generation, which are periods when electricity prices typically spike. Furthermore, due to the significantly larger microgrid capacity in this scenario, the amount of energy sold to the grid increases substantially compared to the original results.

Baseload Guarantee	Energy Purchased (MWh)	Avg. Purchase Price (€/MWh)	Energy Sold (MWh)	Avg. Sale Price (€/MWh)
80%	32 572	39.5	71 728	11.8
90%	26 442	39.8	121 711	13.3
100%	2 403	149.01	1 094 833	17.7

Table 10: Energy purchased and sold on the spot market under different baseload guarantee levels for the 20 MW baseload and 10 MW grid connection case.

Table 11 shows the breakdown of the costs: investment cost, fixed yearly costs, operational costs, and LCOE. Since more energy is sold, the operational costs are lower compared to the original results. The fixed costs however increase a lot, resulting in a rather large increase in the LCOE for this case.

Baseload Guarantee	Operational Cost (M€/year)	Fixed Cost (M€/year)	Total Investment (M€)	LCOE (€/MWh)
80%	2.6	6.0	139	49
90%	0.6	8.4	195	51
100%	-19	44.7	1 031	146

Table 11: Cost comparisons between microgrid configurations.

5.1.3 Stress Testing Microgrid Performance Under Low Production Scenarios

To test the robustness of the configurations, it is useful to assess how they perform under worse conditions than those observed in optimized year, 2024. For this, simulations of wind production, photovoltaic production, and spot prices were run to generate 1000 scenarios. The simulation with the lowest wind and photovoltaic production, paired with the corresponding spot price curve, was used to stress-test the solution. The simulations can be seen in Figure 14.

The mean production efficiency of the 1000 simulations was 27.9% for wind production and 10.2% for photovoltaic production. The standard deviation was 1.7 percentage points for wind and 0.4 percentage points for photovoltaic. In the worst-case scenario, wind production decreased to 22.0% and photovoltaic production to 9.1%. The spot price curves generated to match the wind production had a mean price of 41.5 €/MWh, while in the worst production scenario, the mean spot price rose to 45.6 €/MWh.

The simulation is used to evaluate how the microgrid configurations perform under adverse conditions. A summary of the relevant key performance metrics is presented in Table 12.

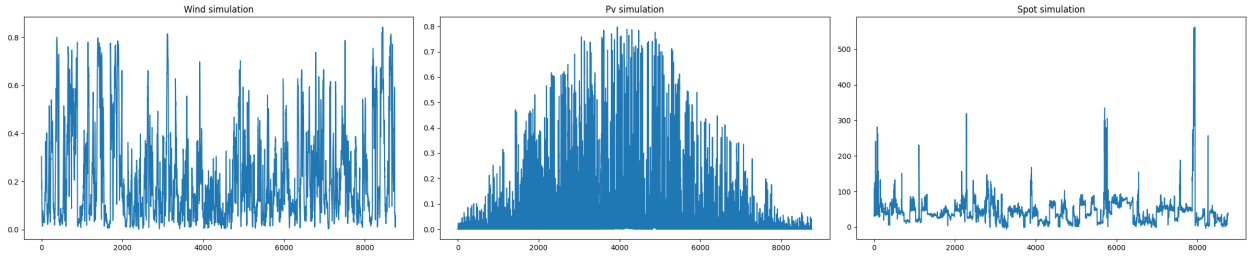

Figure 14: Yearly profiles of the worst-case scenario used for stress testing. The photovoltaic and wind curves show the lowest simulated production levels, with the spot price reflecting the same simulation as the wind data.

Baseload coverage			
	80%	**90%**	**100%**
Baseload coverage	71%	81%	99%
Change	-11%	-9%	-1%
Spotmarket trading			
	80%	**90%**	**100%**
Bought	130 214 MWh	117 172 MWh	49 281 MWh
Price	48.5 €/MWh	49.9 €/MWh	58.0 €/MWh
Sold	0 MWh	423 MWh	105 721 MWh
Price	- €/MWh	26.9 €/MWh	30.0 €/MWh
Operational costs			
	80%	**90%**	**100%**
Operational cost (yearly)	7.0 M€	6.2 M€	-0.30 M€
Change	+2.3 M€	+2.9 M€	+0.24 M€

Table 12: Summary of key performance metrics under different baseload coverage scenarios.

Firstly, all configurations show a decrease in baseload coverage, which is expected given the significantly lower production; any other result would be questionable. Notably, the 100% coverage case still comes close to meeting its target, which is encouraging considering the extremity of the scenario. The 80 and 90% case, however, is more affected and falls to 9 percentage points short of their coverage requirement. Interestingly, the relative drop is greater in the 80% case, since it had previously exceeded its target under normal conditions.

The spot market trading also behaves as expected, with higher volumes on the buy side and lower volumes on the sell side. This is a natural consequence of the reduced production of wind and photovoltaic power, which leads to an increased dependence on external energy purchases.

The average purchasing price increased in all three cases, which is expected given the lower renewable energy generation, leading to higher spot market prices. In the 90% case, the price spread has now aligned with expectations, showing a higher average purchasing price than selling price. The same trend is observed in the 100% case as before; however, the average selling price has also increased. This is likely a result of the overall rise in market prices. Since the spread between the purchasing and selling prices remains relatively stable, the increase in selling price does not indicate a fundamental change in market dynamics but rather reflects a general upward shift in price levels.

Lastly, it is also important to assess how operational costs are affected, as cost efficiency is one of the main objectives of the microgrid. In the low-production scenario, all three baseload coverage cases are negatively impacted, with operational costs increasing by 2.3, 2.9, and 0.24 million euros, respectively. The reduced local production increases the reliance on external electricity imports and the use of the costly diesel generator. Additionally, the higher spot prices in this scenario make external purchases more expensive. Although higher spot prices also enable the microgrid to sell energy at a better price, the net effect on costs remains negative for the three configurations.

Summary and Insights from Stress Testing

The stress test confirms the robustness of the configurations under challenging conditions. While all setups experience reduced performance, they still maintain reasonable levels of reliability. The trade-offs between robustness and cost become particularly apparent, underlining the importance of strategic system sizing.

5.1.4 Worst-Case Optimized Microgrid Design

To evaluate the robustness of the microgrid designs, a "worst-case optimization" was performed, in which the system was designed to meet baseload coverage requirements (80%, 90%, 100%) under a low-production scenario. This approach allows to assess the additional capacity required to fulfil baseload guarantees under highly unfavourable conditions. The optimization algorithm was run for this worst-case scenario, and the resulting microgrid configurations are presented in Table 13.

Baseload Guarantee	Wind Turbines	PV Panels	Batteries	Total Capacity (MW)
80%	5	5	0	40
Change	+2	+2	+0	+16
90%	6	6	1	49
Change	+2	+2	+0	+16
100%	54	0	7	385
Change	+37	-3	-4	+252

Table 13: Installed capacity in different microgrid configurations under varying baseload guarantees in worst-case production scenario.

The new microgrid configurations, optimized for the worst-case scenario, demonstrate notable changes in the optimal design. The most significant change is observed in the 100% baseload guarantee scenario, where the number of wind turbines increased by 37, while the number of photovoltaic panels decreased by 3. This adjustment probably reflects the need for a more stable and reliable energy source during periods of low photovoltaic production. As discussed previously, wind power tends to have less variability compared to photovoltaic power, particularly during critical periods, but it comes at a higher investment cost. It is therefore reasonable to assume that the optimization favours wind in this scenario to ensure continuous baseload coverage. Given the large installed wind capacity, the marginal benefit of additional photovoltaic power becomes less, which could explain the reduction in photovoltaic panels.

Another interesting observation is the low utilization of batteries in all three scenarios. One possible explanation is that batteries involve relatively high costs and energy losses during charging and discharging, making them less attractive unless absolutely necessary to meet baseload coverage targets.

Looking at the total capacity for the three scenarios, there is a dramatic increase in the 100% coverage case, with a total installed capacity of 385 MW, almost 10 times more than in 80%. From a practical point of view, it is highly questionable whether this is a reasonable solution. Significant consideration should be given to whether the economic costs are justified, as this configuration results in a considerably more expensive investment compared to the lower baseload guarantee cases.

5.2 Extension of scope

The energy market, as discussed earlier, is a rather complex system with many layers. As a result, certain aspects have been intentionally excluded from this project, leaving several possible directions for future extensions.

The most significant limitation of the current scope is the focus on only one of the four electricity markets: the day-ahead (spot) market. To broaden the scope, a Power Purchase Agreement (PPA) could be introduced, in which the data center receives a fixed amount of electricity each hour or is supplied with electricity based on the production of the contract counterpart. A potential drawback of the latter approach is that the counterpart is likely to have its production located near the data center (depending on the specific type of contract). To maintain sustainability in electricity supply, the PPA would need to be tied to renewable energy, as is typically the case. In such a scenario, the production patterns of both the microgrid of the data center and the counterpart of the contract would be highly correlated. In other words, the data center would receive electricity at the same hours that it can generate its own power, which could make a PPA contract less advantageous.

Beyond the PPA market, the data center could also operate in the intraday and balancing (regulation) markets. In a realistic setup, the microgrid would make daily predictions to submit bids to the day-ahead market. Since forecasts rarely perfectly match actual out-

come, these predictions could then be updated and acted upon in the intraday market, up to one hour before physical delivery. During delivery hours, further adjustments could be made in the balancing market. This would apply to both energy production and purchasing forecasting.

In this project, the problem setup assumes that whenever production and market orders cannot cover the baseload, a diesel generator must be used to meet the energy demand. However, what has not been discussed is the possibility of curtailing or alternatively shutting down parts of the datacenter, thereby reducing the baseload when coverage is not possible. The reason for this omission is that the specific costs, constraints, and feasibility of such actions would be highly dependent on the particular datacenter in question. Therefore, to maintain a more general approach, this has not been taken into account. However, if this optimization were to be applied in a practical implementation, such a curtailment could be introduced as an additional variable.

In addition, the scope could be extended by introducing different types of penalties for carbon emissions from the diesel generator. In the current setup, the diesel generator is only penalized through its high cost of production.

6 Conclusion

This study investigates how a microgrid can be configured in the most cost-effective way using wind and photovoltaic power together with battery energy storage to provide data centers with a reliable energy supply that covers their baseload demand. To achieve this, a customized cost evaluation function was developed and a particle swarm optimization (PSO) algorithm was used to identify the most cost-effective microgrid configuration. The results show that the optimal combination of wind, solar, and batteries depends on several factors, such as production profile and baseload coverage requirements, but generally favours a higher share of wind power combined with moderate battery capacity and solar power, for the bidding zone considered (SE2). Scenario analyses further show that, under the current assumptions, purchasing the entire baseload from the spot market is more cost-effective than investing in a microgrid. However, this outcome may change shortly due to evolving cost structures and technological advancements. These findings can serve as a guideline for designing cost-efficient and sustainable microgrids, and the developed method is flexible enough to be applied with different input parameters.

References

Achterberg, T. (2009). Scip: solving constraint integer programs. *Mathematical Programming Computation*, 1:1–41.

Bansal, J. C., Singh, P., Saraswat, M., Verma, A., Jadon, S. S., and Abraham, A. (2011). Inertia weight strategies in particle swarm optimization. In *2011 Third world congress on nature and biologically inspired computing*, pages 633–640. IEEE.

Bertail, P. and Dudek, A. E. (2024). Optimal choice of bootstrap block length for periodically correlated time series. *Bernoulli*, 30(3):2521–2545.

Bowden, C. (2023). Understanding power markets: Merit order and marginal pricing. *Squeaky*.

Clerc, M. (2010). *Particle swarm optimization*, volume 93. John Wiley & Sons.

Dudek, A. E. (2018). Block bootstrap for periodic characteristics of periodically correlated time series. *Journal of Nonparametric Statistics*, 30(1):87–124.

Eberhart, R. and Kennedy, J. (1995). Particle swarm optimization. In *Proceedings of the IEEE international conference on neural networks*, volume 4, pages 1942–1948. Citeseer.

Eberhart, R. C. and Shi, Y. (2000). Comparing inertia weights and constriction factors in particle swarm optimization. In *Proceedings of the 2000 congress on evolutionary computation. CEC00 (Cat. No. 00TH8512)*, volume 1, pages 84–88. IEEE.

Ekoue, M. K., Woerman, M., and Clastres, C. (2025). Intermittency and uncertainty in wind and solar energy: Impacts on the french electricity market. *Energy Economics*, 142:108176.

Energimarknadsinspektionen (2023). Elmarknaden.

Energimyndigheten (2024). År 2023 bidrog vindkraften med cirka 21 procent till sveriges elproduktion.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the econometric society*, pages 987–1007.

Gaspar, I., Castro, R., and Sousa, T. (2021). Optimisation and economic feasibility of battery energy storage systems in electricity markets: The iberian market case study. *Journal of Cleaner Production*, 324:129255.

Harrison, R. L. (2010). Introduction to monte carlo simulation. In *AIP conference proceedings*, volume 1204, page 17.

HemSol (2025). Solcellspark: Läs om solparker i sverige de största i världen!

Huang, M.-Y., Huang, L.-Y., Zhong, Y.-X., Yang, H.-W., Chen, X.-M., Huo, W., Wang, J.-Z., Zhang, F., Bai, B., and Shi, L. (2025). Milp acceleration: A survey from perspectives of simplex initialization and learning-based branch and bound. *Journal of the operations research society of china*, 13(1):1–55.

Jakobsson, A. (2021). *An Introduction to Time Series Modeling*. Studentlitteratur, Lund, Sweden, 4 edition.

Jarbratt, G., Jautelat, S., Linder, M., Sparre, E., van de Rijt, A., and Han Wong, Q. (2023). Enabling renewable energy with battery energy storage systems. *McKinsey Company*.

Lindholm, K. (2023). Handel på elbörsen. *Energiföretagen*.

Poli, R., Kennedy, J., and Blackwell, T. (2007). Particle swarm optimization: An overview. *Swarm intelligence*, 1:33–57.

Politis, D. N. and Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical association*, 89(428):1303–1313.

Shi, Y. and Eberhart, R. (1998). A modified particle swarm optimizer. In *1998 IEEE international conference on evolutionary computation proceedings. IEEE world congress on computational intelligence (Cat. No. 98TH8360)*, pages 69–73. Ieee.

Shi, Y. and Eberhart, R. C. (1999). Empirical study of particle swarm optimization. In *Proceedings of the 1999 congress on evolutionary computation-CEC99 (Cat. No. 99TH8406)*, volume 3, pages 1945–1950. IEEE.

SMHI (2025). Solstrålning.

Srivathsan, B., Sorel, M., Sachdeva, P., Bhan, A., Batra, H., Sharma, R., Gupta, R., and Choudhary, S. (2024). Ai power: Expanding data-center capacity to meet growing demand. *McKinsey Company*.

Svensk Solenergi (2025). Introduktion till vår statistik.

Svenska Kraftnät (2024a). Kapacitetsberäkning.

Svenska Kraftnät (2024b). Sveriges elnät.

Svenska Kraftnät (2025). Om elmarknaden.

Thirunavukkarasu, G. S., Seyedmahmoudian, M., Jamei, E., Horan, B., Mekhilef, S., and Stojcevski, A. (2022). Role of optimization techniques in microgrid energy management systems—a review. *Energy Strategy Reviews*, 43:100899.

Trustio AB (2025). Varför är Sverige uppdelat i flera elområden? [Online; accessed 2025-06-04]. Bild över Sveriges elområden använd.

Usaola, J. (2014). Synthesis of hourly wind power series using the moving block bootstrap method. In *2014 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, pages 1–6. IEEE.

Appendix

Table 14: Correlation between wind power production and spot price (2019–2024)

Year	Correlation
2019	-0.116
2020	-0.205
2021	-0.404
2022	-0.429
2023	-0.373
2024	-0.415

Table 15: Correlation between photovoltaic production and spot price (2015–2022)

Year	Correlation
2015	-0.069
2016	0.079
2017	0.134
2018	0.158
2019	-0.061
2020	-0.035
2021	0.045
2022	-0.025

Table 16: Correlation between wind and photovoltaic production (2019–2022)

Year	Correlation
2019	-0.213
2020	-0.307
2021	-0.207
2022	-0.235

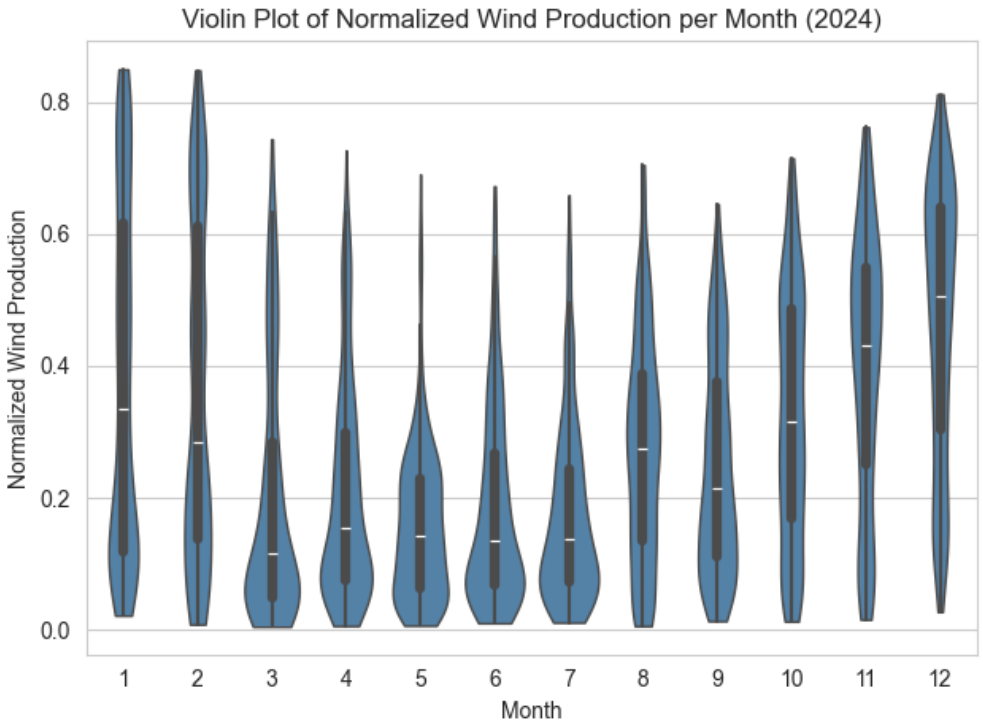

Figure 15: Normalized wind production by month for 2024. The plot shows seasonal variations in both the level and variability of wind output across the year.

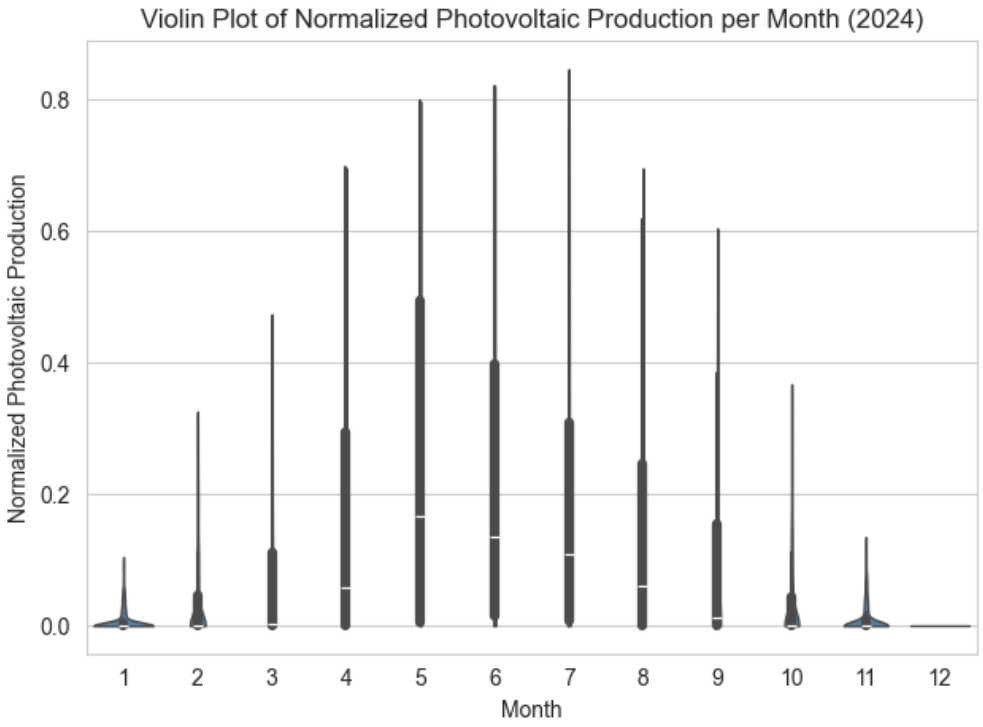

Figure 16: Normalized photovoltaic production by month for 2024.