

---

# **Grid-Aware Sizing and Operation of Distributed Energy Resources**

---

Zur Erlangung des akademischen Grades eines  
**DOKTORS DER INGENIEURWISSENSCHAFTEN**  
*(Dr.-Ing.)*

von der KIT-Fakultät für Informatik  
des Karlsruher Instituts für Technologie (KIT)  
genehmigte

Dissertation  
von

**Simon Grafenhorst**

Tag der mündlichen Prüfung: 15.12.2025  
Hauptreferent: **Prof. Dr. Veit Hagenmeyer**  
Korreferent: **Prof. Dr. Albert Moser**

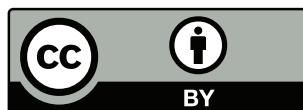
Simon Grafenhorst 

Institute for Automation and Applied Informatics (IAI)  
Karlsruhe Institute of Technology (KIT)

DOI: [10.5445/IR/1000189442](https://doi.org/10.5445/IR/1000189442)

This work is licensed under a Creative Commons Attribution 4.0 International License  
(CC BY 4.0).

<https://creativecommons.org/licenses/by/4.0/>



# Abstract

---

Due to declining costs and policy measures promoting renewable energy sources, the number of renewable energy systems is increasing worldwide. This development demands investments in distribution grid infrastructure and intelligent control mechanisms that compensate fluctuating energy generation from renewable sources. Simultaneously, quantifying these requirements by distribution system operators is impossible without adequate measurement instrumentation. Key components, such as transformers, have long life expectancies and are often not equipped with communication or measurement hardware. As an alternative, smart household appliances are increasingly widespread, and some models are equipped with measurement hardware for monitoring the voltage at the socket outlet. For uninstrumented nodes, pseudo-measurements can be generated using a feature propagation algorithm. This completes voltage monitoring across the entire distribution grid. With knowledge of distribution grid voltage levels and congested lines, Distributed Energy Resources (DERs) can be strategically sized and scheduled to support the grid most economically. This planning and dispatch scheduling process constitutes an optimization problem, which may be solved analytically or heuristically with distinct advantages for either approach. As part of this procedure, power flow calculations are performed iteratively, which is computationally intensive. By leveraging Graph Neural Networks (GNNs), the results of such calculations can be efficiently approximated. The accuracy and efficiency of these approximations are highly dependent on the machine learning model's underlying architecture.

This dissertation investigates the measurement accuracy of widely available smart plugs, demonstrating their sufficient precision for voltage monitoring in distribution grids. A fast feature propagation algorithm supplements these measurements with pseudo-measurements, completing the monitoring across entire distribution grids. Through a comparison of heuristic and analytical optimization approaches for multi-energy DERs, this research establishes that DER components can be accurately described using linear equations in Mixed-Integer Linear Program (MILP) formulations that yield precise analytical solutions and can be computed efficiently. An iterative algorithm adapts the composition and dispatch schedule of a multi-energy DER to the surrounding distribution grid. To accelerate computation, GNNs are employed to infer the algorithm's solutions. A two-stage, pre-trained model that incorporates a first-stage power flow inference model and a second stage that approximates the algorithm's results is compared to a baseline, single-stage GNN that serves as a direct substitute for the iterative algorithm. The pre-trained GNN approximates minimal power requirements for grid congestion mitigation with greater accuracy than the cold-start model in all tested scenarios, including calculations on real-world distribution grid topologies excluded from the training datasets.



# Zusammenfassung

---

Aufgrund sinkender Kosten und politischer Maßnahmen zur Förderung regenerativer Energiequellen steigt die Anzahl erneuerbarer Erzeugungsanlagen weltweit an. Diese Entwicklung erfordert Investitionen in Verteilnetzinfrastruktur und intelligente Steuerungsmechanismen, welche die schwankende Energieerzeugung kompensieren. Gleichzeitig ist die Quantifizierung dieser Erfordernisse durch die Verteilnetzbetreiber ohne ausreichende Messinstrumentierung unmöglich. Infrastrukturkomponenten, wie Transformatoren, haben eine lange Lebensdauer und sind oftmals nicht mit Kommunikations- oder Messhardware ausgestattet. Als Alternative werden Smart Home Geräte mit Hardware zur Spannungsüberwachung an der Steckdose immer populärer. Können in den Netzknoten keine direkten Messungen vorgenommen werden, lassen sich mithilfe eines Feature-Propagation-Algorithmus Pseudo-Messwerte generieren. Damit wird die Spannungsüberwachung über das gesamte Verteilnetz vervollständigt, selbst wenn nicht alle Knoten instrumentiert sind. Mit detaillierten Kenntnissen über Netzengpässe lassen sich dezentrale Energieressourcen DERs auslegen, um das Netz möglichst ökonomisch zu unterstützen. Dieser Planungs- und Steuerungsprozess stellt ein Optimierungsproblem dar, das analytisch oder heuristisch gelöst werden kann, wobei jede Methode eigene Vorteile bietet. Im Rahmen dieses Verfahrens werden iterativ rechenintensive Lastflussberechnungen durchgeführt. Durch den Einsatz von Graph Neural Networks (GNNs) lassen sich diese Berechnungen effizient approximieren. Die Genauigkeit und Effizienz dieser Approximationen hängen von der zugrunde liegenden Architektur des Machine-Learning-Modells ab. In dieser Dissertation wird die Messgenauigkeit kommerzieller Smart-Plugs untersucht, wobei deren hinreichende Präzision für die Spannungsüberwachung in Verteilnetzen nachgewiesen wird. Ein schneller Feature-Propagation-Algorithmus ergänzt diese Messungen durch Pseudo-Messwerte und vervollständigt so die Netzüberwachung. Durch den Vergleich heuristischer und analytischer Optimierungsansätze für sektorenübergreifende DERs wird gezeigt, dass DER-Komponenten mit linearen Gleichungen in Mixed-Integer Linear Program (MILP)-Modellen präzise analytische Lösungen liefern und zudem effizient berechnet werden können. Ein iterativer Algorithmus passt die Zusammensetzung und den Fahrplan einer sektorenübergreifenden DER an das umgebende Verteilnetz an. Um die Rechenzeit zu verkürzen, werden GNNs eingesetzt, um die Ergebnisse des Algorithmus zu approximieren. Ein in zwei Stufen vortrainiertes Modell, das zunächst die Lastflussberechnungen und anschließend die Algorithmusergebnisse approximiert, wird einem Einstufigen-GNN gegenübergestellt, das als direkter Ersatz für den iterativen Algorithmus dient. Das vortrainierte GNN approximiert den minimalen Leistungsbedarf zur Netzengpassbehebung in allen getesteten Szenarien, einschließlich realer Verteilnetz-Topologien, die nicht Teil der Trainingsdaten waren, mit höherer Genauigkeit als das Cold-Start-Modell.



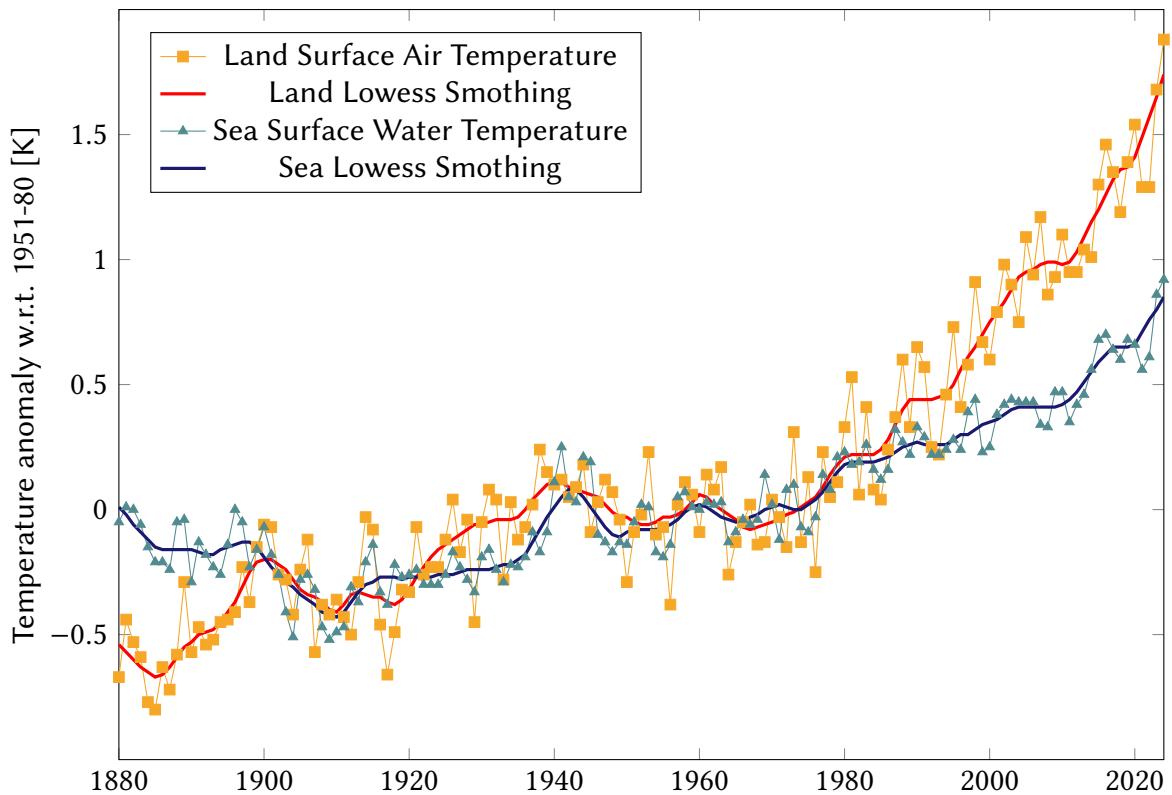
# Contents

---

<b>Abstract</b>	<b>iii</b>
<b>Zusammenfassung</b>	<b>v</b>
<b>Contents</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 Fundamentals</b>	<b>7</b>
2.1 Distributed Energy Resources . . . . .	7
2.2 Scheduling Distributed Energy Resources . . . . .	10
2.2.1 Mixed-Integer Linear Program (MILP) and Mixed-Integer Quadratic Program (MIQP) . . . . .	10
2.3 Power Grids . . . . .	12
2.4 Power Flow Analysis . . . . .	14
2.4.1 Newton-Raphson Solver . . . . .	16
2.4.2 Pandapower . . . . .	18
2.5 Power Grid Graph Representation . . . . .	19
2.5.1 Graph Neural Networks . . . . .	19
2.5.2 Multilayer Perceptrons . . . . .	22
2.5.3 Training Procedure . . . . .	22
<b>3 Distribution Grid Monitoring</b>	<b>25</b>
3.1 Introduction . . . . .	25
3.2 Related Work . . . . .	31
3.3 Method . . . . .	35
3.3.1 Measurement Hardware . . . . .	35
3.3.2 Measurement Data Acquisition . . . . .	35
3.3.3 Data Processing . . . . .	38
3.3.4 Feature Propagation . . . . .	38
3.4 Evaluation . . . . .	42
3.4.1 Evaluating the Smart Plugs Measurement Accuracy . . . . .	42
3.4.2 Distribution Grid Monitoring Error . . . . .	44
3.5 Discussion . . . . .	47
3.6 Conclusion and Outlook . . . . .	49
<b>4 Optimization of Multi-Energy DERs</b>	<b>51</b>
4.1 Introduction . . . . .	51

4.2	Related Work . . . . .	54
4.3	Method . . . . .	57
4.3.1	Evolutionary Algorithms . . . . .	57
4.3.2	Analytical Model . . . . .	57
4.4	Evaluation . . . . .	61
4.4.1	Optimization Results . . . . .	61
4.4.2	Runtime . . . . .	65
4.4.3	Configurability . . . . .	65
4.5	Discussion . . . . .	66
4.6	Conclusion and Outlook . . . . .	68
<b>5</b>	<b>Grid Aware DER Optimization</b>	<b>69</b>
5.1	Introduction . . . . .	69
5.2	Related Work . . . . .	71
5.3	Method . . . . .	73
5.4	Case Studies . . . . .	77
5.4.1	Case Study 1: Proof of Concept . . . . .	79
5.4.2	Case Study 2: Large Energy Systems . . . . .	82
5.5	Discussion . . . . .	85
5.6	Conclusion and Outlook . . . . .	87
<b>6</b>	<b>Graph Neural Networks for Grid-Aware DER Sizing</b>	<b>89</b>
6.1	Introduction . . . . .	89
6.2	Related Work . . . . .	91
6.3	Method . . . . .	93
6.3.1	Training Data . . . . .	93
6.3.2	Power Flow Model . . . . .	94
6.3.3	Model Architecture for Mitigating Congestion . . . . .	97
6.4	Evaluation . . . . .	100
6.4.1	Power Flow Model . . . . .	101
6.4.2	Training the Pre-Trained and Cold-Start Models . . . . .	102
6.4.3	Performance Evaluation of the Pre-Trained and Cold-Start Models	104
6.5	Discussion . . . . .	106
6.6	Conclusion and Outlook . . . . .	108
<b>7</b>	<b>Summary</b>	<b>111</b>
<b>8</b>	<b>Appendix</b>	<b>115</b>
	<b>Bibliography</b>	<b>121</b>

Electricity grids around the world are undergoing significant transformations, driven by a conjunction of economic and environmental factors. The primary environmental driver is the urgent need to stop climate change, which is principally caused by greenhouse gas emissions from burning fossil fuels [Cal+23]. This human-caused dramatic rise in average global surface temperature, both over land and oceans, is illustrated in Figure 1.1. In response, international frameworks such as the International Renewable Energy Agency (IRENA) 1.5° scenario outline a clear pathway to limit global warming by fundamentally transforming energy systems, a transition made feasible by compelling economic incentives [Int+24].



**Figure 1.1:** “Annual (thin lines) and five-year [averaged] (thick lines) [...] temperature anomalies (vs. 1951-1980) averaged over the Earth’s land area and sea surface temperature anomalies (vs. 1951-1980) averaged over the part of the ocean that is free of ice at all times (open ocean).” [NG25] Source: [https://data.giss.nasa.gov/gistemp/graphs\\_v4/](https://data.giss.nasa.gov/gistemp/graphs_v4/)

The falling cost of renewable energy generation is a pivotal economic factor accelerating the energy transition. On a global scale, renewable energy is now the cheapest option for increasing electricity generation in almost every country [Int+24]. In Germany, even low-voltage Photovoltaic (PV) systems connected to the distribution grid generate electricity at a lower cost than lignite, coal, or nuclear power plants [Kos+24]. However, realizing the goals of the energy transition is not without challenges. The increase in renewable generation must be matched by a significant enhancement in the energy transmission capabilities of both transmission and distribution grids to ensure system reliability and stability [Int+24].

Because of the rising cost of carbon emissions [Wor24] and fuel prices [Nev+17], heating with gas and oil is becoming more expensive. This makes electric heating more attractive and leads to more heat pump installations, which changes the characteristic load curve in electric grids. Governmental interventions reinforce this trend by introducing incentives to promote the installation of heat pumps [Bf23]. As electric heating and cooling become more popular, electricity consumption patterns become increasingly dependent on weather conditions. Similarly, the cost of operating fossil fuel-powered vehicles is increasing, while Electric Vehicles (EVs) are becoming more affordable. This necessitates the integration of EV chargers into the existing grid infrastructure. Incorporating such high-power loads of up to several hundreds of kW into the distribution grids poses significant challenges [Ham+24]. On the other hand, renewable energy generation is linked to meteorological conditions and is not necessarily available at times of high demand. PV systems, for example, produce the most power at midday during summer, when there is the most sunlight. This contrasts with the consistent output characteristic of fossil fuel-based generation systems. The intermittent nature of renewable energy production leads to fluctuations in the utilization of the distribution grid. Especially at times with high simultaneous generation, the distribution grid voltage increases due to the high feed-in power. During these times, energy may flow from the low-voltage grids into the medium-voltage grids, increasing the voltage levels in the medium-voltage grids [Mit+15] and the utilization of transformers and lines. Furthermore, if the voltage magnitude is too high, feed-in must be curtailed.

Curtailing decentralized feed-in is associated with costs for the grid operators (§ 13a of the German Energiewirtschaftsgesetz (EnWG)) that are passed on to the consumers. Ideally, the grid operator should be able to detect violations of the permissible voltage range beforehand by extensively monitoring the distribution grid. This enables the estimation of how and when the distribution grid needs to be upgraded. However, grid operators often lack this information about the current grid state. This information deficit could potentially be addressed by leveraging other data sources. For instance, the adoption of smart home technology is experiencing significant growth [Mar22] and may offer a viable means for enhancing grid observability.

Within this expanding market, several kinds of devices incorporate measurement hardware specifically designed to quantify the energy consumption of the device itself or connected appliances. This functionality is achieved by measuring voltage and current concurrently. The data acquired via these measurements holds potential utility for enhancing the monitoring of the electrical distribution grid. This motivates the first research question of this thesis: **To what extent can smart plugs be used for distribution grid monitoring?**

To realize complete monitoring of the distribution grid with the help of this measurement data, the voltage values at each node must be known. If nodes are not instrumented, the measured values must be calculated or derived differently. Rossi et al. [Ros+21] apply the principle of heat propagation to information propagation in graphs and present a feature propagation algorithm. The propagation of heat can be calculated using the heat diffusion equation as a function of thermal conductivity and heat sources. A direct analogy exists within graph theory: the weights assigned to graph edges can be conceptualized as analogous to thermal conductivity, while nodes acting as origins of information parallel the role of heat sources. This leads to the second research question: **How viable is the use of a feature propagation algorithm for pseudo-measurement generation in distribution grids compared to other related methods?**

If the state of the distribution grid is known, congestion can be mitigated, for example, by using DERs. These decentralized systems can store different forms of energy or convert them into one another, for instance, via combustion engines, electrolyzers, and fuel cells. By combining multiple components into a single DER, operational synergies can be utilized [Gei+07]. The planning and sizing of a multi-energy DER presents a complex optimization problem that can be separated into two fundamental tasks: system composition and operational scheduling. Both optimization problems require a multi-criteria approach in order to find the best possible trade-off between the different goals a DER can pursue. These goals include supporting the distribution grid, minimizing environmental impacts such as carbon emissions, and maximizing economic value through revenue generation. The optimal DER implementation therefore represents a compromise among these objectives, both in the configuration of the system and in its dynamic operational strategy. Depending on the modeling and optimization method, the optimization goal is achieved with varying degrees of quality and computational effort. In order to choose a suitable approach for optimizing a multi-energy DER, another research question is: **What are the advantages and disadvantages of optimizing a multi-energy DER using an evolutionary algorithm, an MILP, and an MIQP?**

When planning and sizing such a multi-energy DER, the surrounding distribution grid should be considered. Depending on the loads and generation in the distribution grid, overloads occur, which should be compensated by the DER. At the same time, the DER should not be planned to be too large and should be as economical as possible.

This problem raises a further research question: **How does the annual cost of a DER designed to provide grid supporting services compare to the annual cost of the most economical design?**

The algorithm developed for planning a DER located in a specific distribution grid area is based on a large number of power flow calculations as a part of an iterative, computationally demanding algorithm. Speedups for power flow calculations can be achieved, for example, by implementing a power flow inferring GNN. It therefore stands to reason that a GNN can also learn the results of the iterative algorithm. However, the two problems of power flow prediction using GNNs and optimizing the DER's power can also be considered separately. Thus, the power flow approximating GNN can be trained using large training datasets and provide the best possible basis for inferring the optimal DER power. This leads to the next research question: **How effective is a pre-trained GNN compared to a cold-start GNN for inferring the optimal DER power?**

In summary, the following research questions are devised:

1. **To what extent can smart plugs be used for distribution grid monitoring?**
2. **How viable is the use of a feature propagation algorithm for pseudo-measurement generation in distribution grids compared to other related methods?**
3. **What are the advantages and disadvantages of optimizing a multi-energy DER using an evolutionary algorithm, an MILP, and an MIQP?**
4. **How does the annual cost of a DER designed to provide grid supporting services compare to the annual cost of the most economical design?**
5. **How effective is a pre-trained GNN compared to a cold-start GNN for inferring the optimal DER power?**

This dissertation is structured as follows: The first chapter outlines the fundamentals, which include DERs, power grids, the grid infrastructure, and power flow calculations. Furthermore, it provides an exposition on GNNs. The second chapter investigates the domain of distribution grid monitoring. It introduces a methodology for monitoring the low-voltage grid that employs smart plugs. This approach is subsequently extended to achieve comprehensive grid observability through the use of pseudo-measurements, which are synthesized for all grid non-instrumented nodes via a feature propagation algorithm. The third chapter addresses the optimization of DERs containing multi-energy components, and presents a comparative analysis of heuristic Evolutionary Algorithms (EAs) and analytic mixed-integer problem formulations. The fourth chapter then builds upon this foundation. It enhances the analytic optimization algorithm by integrating constraints that stem from the surrounding distribution grid. As the monitoring process makes potential congestion scenarios evident, the enhanced algorithm optimizes the DER's operation specifically to mitigate these events.

The primary drawback of the bespoke DER optimization algorithm is its computational complexity. Consequently, the fifth chapter focuses on the development of GNNs to serve as computationally efficient approximators of the algorithm's results. A baseline model is first established using a single-stage GNN approach. In this configuration, the GNN directly infers the optimal DER power required for congestion mitigation from the grid's load and generation components. A second, more advanced methodology is also developed, employing a two-stage GNN architecture. In the first stage, the network is tasked with approximating the power flows within the grid. The resulting power flow estimations are then supplied as input to a second GNN, a design which facilitates the generation of more accurate inference results. The sixth chapter summarizes and critically discusses the findings, encompassing the complete, end-to-end workflow presented throughout the thesis. Finally, the seventh chapter concludes the dissertation, offering a final synthesis of the work and outlining promising directions for future research.



## 2.1 Distributed Energy Resources

Distributed Energy Resources (DERs) are characterized as small (kW-scale) energy generation, controllable loads, or storage systems located at the distribution level of the power grid, contrasting with traditional centralized generation [Hat+07]. DERs intrinsically enable bidirectional energy exchange between nodes within the distribution grid. The capability of energy production, for example, through local generation mechanisms, such as PV systems installed on rooftops and balconies, precipitates a transition in the role of end-users, from mere consumers to “**prosumers**”. However, the definition of DERs also includes certain energy-consuming appliances, particularly those with Demand-Side Management (DSM) capabilities. EV chargers can sometimes be remotely managed by Distribution System Operators (DSOs) to avoid congestion [HAI23]. Regulatory frameworks, such as § 14a of the German EnWG, further extend the number of manageable appliances by enabling DSOs to control high-load systems like heat pumps. This regulatory approach underscores the growing importance of integrating controllable devices with DSM to meet multiple needs, also called active DSM [Di +18]. This control mechanism allows different objectives to be combined, such as minimizing costs, maximizing self-consumption, and grid support [Mat+12].

Although DERs can be individual components, a combination of components can also be classified as a DER. This integrated approach enables the provision of multiple forms of energy, such as electricity, thermal energy, and hydrogen as a gaseous energy carrier. Henceforth, DERs utilizing multiple forms of energy and comprising multiple individual components are termed multi-energy DERs. An example of such a multi-energy DER is the energy hub [Gei+07]. Energy flows within an energy hub can be effectively modeled using a coupling matrix that quantifies the relationships between energy inputs and outputs. The relationship between the inputs and outputs of the energy hub is defined by

$$\underbrace{\begin{pmatrix} L_1 \\ L_2 \\ \vdots \\ L_n \end{pmatrix}}_{Output} = \underbrace{\begin{bmatrix} C_{1,1} & C_{2,1} & \dots & C_{n,1} \\ C_{1,2} & C_{2,2} & \dots & C_{n,2} \\ \vdots & \vdots & \ddots & \vdots \\ C_{1,n} & C_{2,n} & \dots & C_{n,n} \end{bmatrix}}_{Coupling} \times \underbrace{\begin{pmatrix} P_1 \\ P_2 \\ \vdots \\ P_n \end{pmatrix}}_{Input}. \quad (2.1)$$

The coupling matrix provides a mathematical representation of the energy hub, where each element  $C_{m,n}$  defines the conversion factor from input energy carrier  $m$  to output carrier  $n$ . The diagonal coefficients  $C_{n,n}$  correspond to instances of direct pass-through where the energy carrier remains unchanged. Conversely, the off-diagonal elements  $C_{m,n}$  for  $m \neq n$  model the transformation between two distinct energy carriers. The magnitude of these off-diagonal coefficients is physically significant: values between 0 and 1 represent energetically lossy conversions, while values exceeding 1 denote processes of energy multiplication, such as the thermal output of a heat pump relative to its electrical input. A coefficient of  $C_{m,n} = 0$  signifies that no technological pathway exists within the hub to convert energy carrier  $m$  to  $n$ .

The model is extended beyond this input-output mapping to incorporate energy storage systems, which necessitates the inclusion of their dynamic behavior. Consequently, the complete energy hub formulation integrates the storage dynamics:

$$\underbrace{\begin{pmatrix} M_1 \\ M_2 \\ \vdots \\ M_n \end{pmatrix}}_{\text{Storage input}} = \underbrace{\begin{bmatrix} S_{1,1} & S_{2,1} & \dots & S_{n,1} \\ S_{1,2} & S_{2,2} & \dots & S_{n,2} \\ \vdots & \vdots & \ddots & \vdots \\ S_{1,n} & S_{2,n} & \dots & S_{n,n} \end{bmatrix}}_{\text{Storage coupling}} \times \underbrace{\begin{pmatrix} \dot{E}_1 \\ \dot{E}_2 \\ \vdots \\ \dot{E}_n \end{pmatrix}}_{\text{Energy flow}}. \quad (2.2)$$

The combination of Equation (2.1) and Equation (2.2) yields the complete system formulation:

$$\mathbf{L} = \mathbf{CP} - \mathbf{S}\dot{\mathbf{E}}. \quad (2.3)$$

In order to operate an energy hub with minimal costs, the individual components must be planned in the most economical way. By linearizing all conversion processes and constraints, an MILP can be used for optimization. Due to the combination of multiple components and interactions with the connected energy grids, the design and planning of energy hubs is complex and multifaceted.

The energy hub proposed by Geidl et al. [Gei+07] represents a specific kind of multi-energy DER that utilizes synergy effects among multiple energy carriers, such as electricity and natural gas. Their work provides models and methods to optimize the power exchange between the carriers, thereby minimizing costs or emissions, or maximizing efficiency, availability, and security.

Likewise, a multi-energy DER can be configured to be controlled by the DSO to provide grid support services. Nerowski et al. [Ner+25], for instance, develop a grid congestion mitigation strategy that demonstrates this feature. In this scheme, the energy hub exchanges information with the grid operator, communicating its optimal operating costs. The DSO then responds with load targets for specific time stamps, prompting the energy hub to re-optimize its operation to meet these requirements. This process, however, leads to higher

operational costs for the energy hub, which allows a monetary value to be assigned to the requested grid-supportive service. Through the local integration of energy generation, storage, and consumption, the dependence on long-haul energy transport is minimized. This can postpone or limit the financial outlay for grid expansion projects, leading to a concurrent reduction in energy losses and an increase in overall system efficiency [Ela+23; Zho+22].

Furthermore, energy hubs enable further integration of Renewable Energy Sources (RESs) into the distribution grids. By providing localized storage and flexible demand, they assist in avoiding renewable energy curtailment, encouraging greater penetration of RESs into the grid [Liu+12]. The operational flexibility of an energy hub allows it to attenuate peaks in both energy demand and intermittent renewable generation, contributing to a more efficient grid operation.

Energy hubs also bolster distribution grid resilience [Wan+17]. As localized DERs, they can supply backup power during infrastructure failures, maintaining supply to critical loads. This function enhances the overarching reliability and security of the energy system [Wan+17] by providing substantial operational flexibility, primarily achieved through inter-sectoral energy exchange [Eid+16].

The economic benefit of energy hubs comes from the demand and generation shifting to periods of more favorable pricing, thereby reducing overall energy costs [Wan+17]. For instance, during intervals of high electricity prices, natural gas can be used for on-site power generation. Conversely, when electricity prices are low, excess electrical energy can be converted, for example, to produce synthetic gas for potential injection into the gas network.

## 2.2 Scheduling Distributed Energy Resources

The scheduling of a multi-energy DER constitutes a complex optimization problem. It involves orchestrating the various components with their respective energy inputs and outputs, subject to a set of constraints. These constraints include operational limits, the size of individual components, and external requirements. For example, the DSO could request the DER to provide a specific load or generation profile to support grid operations. Furthermore, weather conditions govern the efficiency of the renewable generation components.

The objective function, therefore, is multifaceted, designed to capture all the potential benefits a DER can provide, such as direct monetary savings, revenue from grid services, and penalties for deviating from a target power curve. The nature of the optimization problem changes significantly depending on the primary objective. When the sole objective is to follow a specified power target, the problem simplifies into a disaggregation task, where the central challenge is to allocate the required power among the available components. Conversely, if the objective is purely to minimize internal costs or maximize revenue, the components each operate at their own cost-optimal setting without any coordinated interaction for broader system goals. Therefore, optimal scheduling is achieved by balancing these extremes, finding a sophisticated control strategy that respects all constraints while intelligently coordinating all components to achieve the best possible outcome across a range of competing objectives.

### 2.2.1 Mixed-Integer Linear Program (MILP) and Mixed-Integer Quadratic Program (MIQP)

A mathematical optimization problem consists of an objective function and a search space. The search space is formulated as a set  $A$  whose elements satisfy specific constraints. These points represent all valid solutions to the problem. The objective function maps these elements to a value that is to be minimized or maximized. Without loss of generality, a minimization problem can be formulated as

$$f(\mathbf{x}_m) \geq f(\mathbf{x}) \Leftrightarrow -f(\mathbf{x}_m) \leq -f(\mathbf{x}) \quad \forall \mathbf{x} \in A. \quad (2.4)$$

A Linear Program (LP) is a type of mathematical optimization in which the objective is to find a vector  $\mathbf{x}$  that minimizes an objective function,  $f(\mathbf{x}) = \mathbf{c}^T \mathbf{x}$  with  $\mathbf{c} \in \mathbb{R}^n$ , over a search space defined by the convex hull of a set of points. This set of points can be described with  $A\mathbf{x} \leq \mathbf{b}$ ,  $\mathbf{b} \in \mathbb{R}^m$  and  $\mathbf{x} \geq 0$  [BN01]. The  $m \times n$  matrix  $A$  contains the coefficients of all constraints, and the vector  $\mathbf{b}$  is a given right-hand side of the constraints.

Both MILP and MIQP are specific formulations of such problems. An MILP extends the LP approach to include integer and binary decision variables in addition to the continuous coefficients. While an MILP handles linear objective functions, an MIQP is necessary to optimize problems that are described by quadratic functions. In this work, the DER com-

ponents are approximated and implemented with linear models. The error metric used to follow a target from the DSO as closely as possible is a Mean Squared Error (MSE) function. Since this error metric contains a quadratic term, an MIQP formulation is required. Several frameworks exist to formulate such optimization problems. This work utilizes the Pyomo (“a Python-based tool for modeling and solving optimization problems”) [Byn+21; HWW11], and GBOML (“Graph-Based Optimization Modeling Language”) [Mif+22] frameworks, which both offer interfaces to a range of solvers. These include commercial solvers, such as Gurobi [Gur23], as well as open-source solvers, such as SCIP [Bol+24]. The Gurobi solver is employed for all MIQP and MIQP optimization problems in this research, due to its availability for academic purposes at no cost and its high performance.

A distinguishing feature of Pyomo is its fundamentally object-oriented modeling approach. This paradigm allows for the decomposition of a complex optimization problem into its individual components, which are defined as reusable classes. This modular structure provides an intuitive and maintainable modeling environment within object-oriented programming languages. The Pyomo framework supports a broad spectrum of problem formulations, including linear programs, quadratic programs, non-linear programs, and MILP and MIQP. This flexibility makes it well-suited for the possibly evolving models of the DER formulation central to this work.

In contrast, GBOML is a more recent framework for graph-based DER modeling. Although limited to MILP formulations, its design facilitates rapid model development. To ensure the future extensibility and reusability of the models developed within this thesis, GBOML is only employed for the initial, linear DER model.

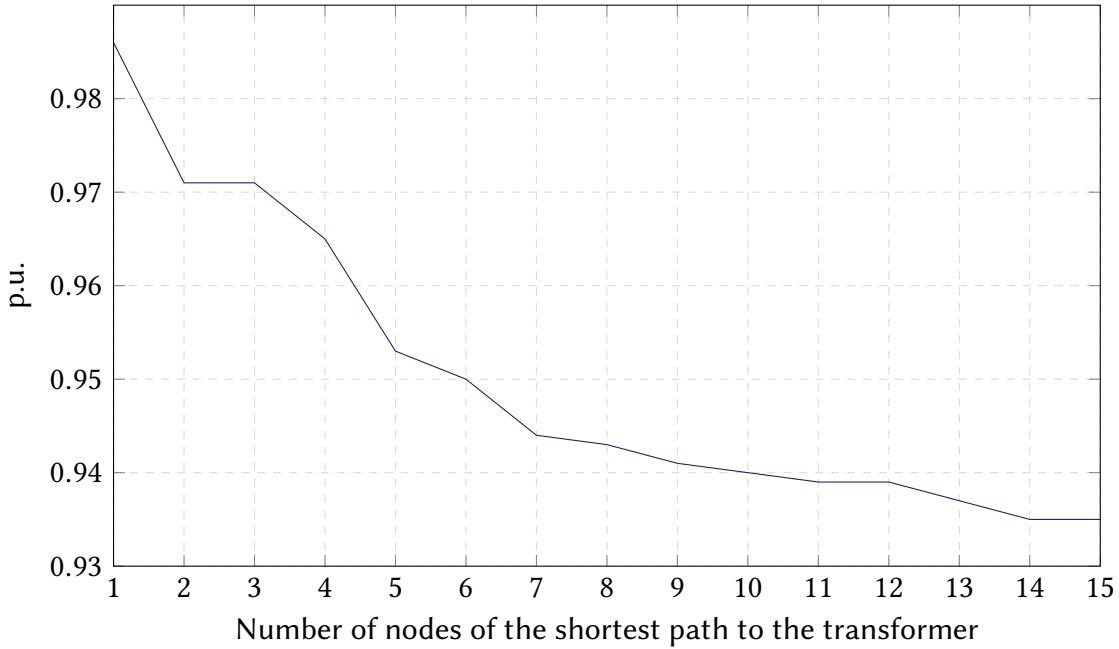
## 2.3 Power Grids

Power grids can be characterized based on their voltage level. In Germany, the distribution grids usually operate at the low voltage level of 400 V and the medium voltage level between 6 kV and 60 kV, while the high voltage level between 60 kV and 220 kV is mostly found in regional transmission grids. In order to cover long distances, transmission grids with voltages between 220 kV and 380 kV and Direct Current (DC) high-voltage grids with voltages up to 525 kV are used [Ger25].

Electrical distribution grids are predominantly designed with a three-phase configuration, which enables the delivery of multiple voltage levels. A typical example is the low-voltage distribution grid, where nominal voltages of 230 V and 400 V are available to consumers. These voltage levels are achieved through distinct connection methodologies: 230 V is obtained by connecting to a phase and neutral conductor, whereas 400 V is achieved through phase-to-phase connections. However, power flow analysis is often conducted for a single phase [LLW17; Wid+10].

Modern distribution grids contain DERs, measurement devices, transformers, switching infrastructure, and more. Increasing penetration of distributed generation and demanding loads, such as EV chargers, must go hand in hand with adjustments to the grid infrastructure. Smart meters enable time-dependent energy prices, enabling the consumer to react to price signals. With this method, incentives can be created such that the consumer voluntarily supports the grid. The smart meters track the energy consumption in configurable time intervals. In the European Union, a 15-minute time interval was deemed appropriate in 2012 [Swa+21]. Therefore, new energy prices can be set for every 15-minute time window.

Using smart meters, load and generation in distribution grids can be shifted through consumers' reactions to price signals. However, the price signal for household consumers is currently not affected by grid congestion but rather the energy prices at the exchanges. However, the distribution system operator is able to use other tools to enhance the capabilities of the distribution grid. A regulated distribution transformer's ratio between the primary and secondary sides can be tuned in such a way that the voltage levels in the distribution grid are within permissible bounds. Therefore, even if voltage levels on the primary high voltage side increase, for example, due to high renewable generation in the area, voltage levels on the secondary low voltage side can remain within permissible bounds. Similarly, when the level of distributed generation in the low-voltage grid is high, the voltage at the transformer's low-voltage side can be reduced, and the voltages stay within bounds at every node in the grid. Mitra et al. [Mit+15] exploit the inherent correlation between voltage magnitudes at low-voltage and medium-voltage levels within the distribution system to derive a connectivity model.



**Figure 2.1:** Illustration of the voltage drop in the distribution grid depending on the length of the shortest path to the transformer.

In distribution grids, line impedance causes voltage drops that become more severe with electrical distance. Consequently, grid participants connected to distant buses may experience low voltage levels. Furthermore, high demand throughout the grid leads to large power transmission, particularly through lines near the feeder. This results in elevated line temperatures and thermal stress. These grid congestion issues are identified using power flow analysis, particularly steady-state power flow calculations, which solve the nodal power balance equations. For such nonlinear systems, the Newton-Raphson method is a standard iterative solver. Figure 2.1 depicts an exemplary 400 V European radial feeder with 20 kW loads at all nodes. The resulting nodal voltages, shown in Figure 2.1, can be calculated using the Newton-Raphson algorithm. All power flow analyses in this thesis are performed using the Pandapower library. Voltages are reported in per-unit (per unit (p.u.)), defined as the ratio of actual to nominal voltage, and are evaluated against regulatory limits, such as the 0.9 to 1.1 p.u. range specified by the EN-50160 standard.

## 2.4 Power Flow Analysis

A power flow analysis constitutes a fundamental numerical analysis of an interconnected electrical grid. Its primary purpose is to determine the steady-state operating conditions of the system. The analysis is based on four fundamental Alternating Current (AC) power grid parameters for each bus: the voltage magnitude  $|V|$ , the voltage angle  $\theta$ , the real power  $P$ , and the reactive power  $Q$  [Glo+22]. To manage the inherent complexity of power grids, these studies employ simplified representations. Key aspects of the simplification include representing a balanced three-phase system with a single line and the p.u. system, which normalizes voltage values. Critically, this analysis focuses solely on the system's behavior under normal, steady-state conditions. This provides an overview of the grid's state, excluding transient phenomena such as faults or switching events.

The power-flow analyses conducted within this thesis operate under the unifying assumption that for every bus in the network, the real power and reactive power are specified quantities. Consequently, the voltage magnitude ( $|V|$ ) and phase angle ( $\theta$ ) are treated as the unknown variables to be determined for all buses. These buses are referred to as *PQ* buses.

This contrasts with another common approach that also incorporates *PV* buses for generators, at which the voltage magnitude is a known, controlled setpoint. In order to ensure a valid solution for the power-flow equations, it is essential to designate one bus as the slack bus, also known as the swing or reference bus. The role of this bus is paramount in that it addresses the disparity between the scheduled power generation and the aggregate system demand, encompassing both consumer loads and the inherent network losses. Given that the system's power losses remain unknown until the power-flow calculation reaches convergence, the slack bus functions as the balancing component, thereby supplying or absorbing the required real and reactive power to maintain power equilibrium across the network. In the grid topologies examined in this thesis, the primary transformer, which functions as the connection to the upstream transmission system, is designated as the slack bus.

Solving the power-flow problem involves determining the voltage magnitude, voltage angle, real power, and reactive power for every bus in the network. The basic relationships between these variables is described using the power flow equations, which are derived in the subsequent section based on the work of Glover et al. [Glo+22].

The injected power at bus  $k$  in a three-phase electricity grid is denoted by

$$\begin{aligned} S_k &= P_k + jQ_k \\ &= \sqrt{3}V_kI_k^*. \end{aligned} \tag{2.5}$$

This complex quantity is constituted by its real and imaginary parts, which correspond to the active power ( $P_k$ ) and reactive power ( $Q_k$ ), respectively. The asterisk is used to indicate complex conjugate quantities and  $j = \sqrt{-1}$ .

The electrical characteristics of a transmission line are modeled by its impedance ( $Z$ ). This complex quantity is constituted by a real component, the resistance ( $R$ ), which accounts for real power losses through dissipation, and an imaginary component, the reactance ( $X$ ), which represents the lossless energy storage characteristic of the line in an AC system. The relationship for a line between the buses  $k$  and  $n$  is given by

$$Z_{kn} = R_{kn} + jX_{kn}. \quad (2.6)$$

Whereas the impedance quantifies a circuit's opposition to current, the admittance ( $Y$ ) measures the ease with which a circuit allows current to pass for a given applied voltage. Mathematically, the admittance is simply the reciprocal of the impedance ( $Y = 1/Z$ ). Just as the impedance is deconstructed into the resistance and the reactance, the admittance can be broken down into the conductance  $G$  and the susceptance  $B$ . With the phase angle  $\angle\theta_{kn}$ , the total admittance  $\mathcal{Y}$  of the line between bus  $k$  and bus  $n$  is thus

$$\begin{aligned} \mathcal{Y}_{kn} &= G_{kn} + jB_{kn} \\ &= |\mathcal{Y}_{kn}| \angle\theta_{kn} \\ &= |\mathcal{Y}_{kn}| e^{j\theta_{kn}} \\ &= |\mathcal{Y}_{kn}| (\cos(\theta_{kn}) + j \sin(\theta_{kn})). \end{aligned} \quad (2.7)$$

The currents that are injected at each node can be represented analogously to Ohm's law ( $I = \frac{U}{R}$ ):

$$\begin{aligned} I &= \frac{V}{Z} \\ &= YV. \end{aligned} \quad (2.8)$$

With the terms  $Y_{kn}$  representing elements of the admittance matrix, the injected current at bus  $k$  can be written as

$$I_k = \sum_{n=1}^N Y_{kn} V_n. \quad (2.9)$$

To calculate the complex power injected at bus  $k$ ,  $S_k$ , the Equation (2.9) can be inserted into Equation (2.5):

$$\begin{aligned}
 S_k &= P_k + jQ_k \\
 &= V_k I_k^* \\
 &= V_k \sum_{n=1}^N Y_{kn}^* V_n^*.
 \end{aligned} \tag{2.10}$$

Expanding  $Y_{kn}^*$  and using phasors as in Equation (2.7), the Equation (2.10) becomes

$$\begin{aligned}
 S_k &= |V_k| \sum_{n=1}^N |Y_{kn}| |V_n| e^{j(\delta_k - \delta_n - \theta_{kn})} \\
 &= |V_k| \sum_{n=1}^N |V_n| (\cos(\delta_k - \delta_n) - j \sin(\delta_k - \delta_n)) (G_{kn} - jB_{kn}).
 \end{aligned} \tag{2.11}$$

In the above Equation (2.11), the terms  $\delta_k$  and  $\delta_n$  denote the phase angles at nodes  $k$  and  $n$ , respectively. From Equation (2.11), the power flow equations can be devised as

$$P_k = |V_k| \sum_{n=1}^N |V_n| (G_{kn} \cos(\delta_k - \delta_n) + B_{kn} \sin(\delta_k - \delta_n)), \tag{2.12a}$$

$$Q_k = |V_k| \sum_{n=1}^N |V_n| (G_{kn} \sin(\delta_k - \delta_n) - B_{kn} \cos(\delta_k - \delta_n)). \tag{2.12b}$$

### 2.4.1 Newton-Raphson Solver

The Newton-Raphson method is one of the most popular techniques for solving load flow equations. It exhibits quadratic convergence properties, reaching a solution in significantly fewer iterations than the Gauss-Seidel method. The computational effort in each iteration is, however, higher [Jay16]. The algorithm is based on the linearization of the nonlinear power flow equations and solves them iteratively using Taylor expansion approximation as

$$y \approx f(x_0) + f'(x_0)(x - x_0). \tag{2.13}$$

Letting the approximation be an equality results in the following Equation (2.14) at the  $i$ -th iteration:

$$\begin{aligned} y &= f(x_i) + f'(x_i)(x_{i+1} - x_i) \\ f'(x_i)(x_{i+1} - x_i) &= y - f(x_i) \\ x_{i+1} &= x_i + \frac{y - f(x_i)}{f'(x_i)}. \end{aligned} \quad (2.14)$$

In this Equation (2.14), the first-order partial derivatives of the power flow Equations (2.12) are the Jacobian matrix  $f'(x_i) = J_i$ . The Jacobian matrix can be written as

$$J = \begin{bmatrix} \frac{\partial P}{\partial \delta} & \frac{\partial P}{\partial V} \\ \hline \frac{\partial Q}{\partial \delta} & \frac{\partial Q}{\partial V} \end{bmatrix}. \quad (2.15)$$

The vector  $f(x_i)$  in Equation (2.14) denotes the following column vector (2.16) that represents the differences between specified and calculated power injections. Because the first bus is the slack bus, indices start with 2.

$$f(x) = \begin{bmatrix} P_2^{\text{spec}} - P_2^{\text{calc}} \\ \vdots \\ P_k^{\text{spec}} - P_k^{\text{calc}} \\ Q_2^{\text{spec}} - Q_2^{\text{calc}} \\ \vdots \\ Q_k^{\text{spec}} - Q_k^{\text{calc}} \end{bmatrix} \quad (2.16)$$

The iterative algorithm stops when all elements of this vector  $f(x)$  reach a specified threshold, for example

$$\max|f(x)| < \epsilon. \quad (2.17)$$

### 2.4.2 Pandapower

For the power system analyses in this thesis, the Pandapower library is selected [Thu+18]. The choice of Pandapower is motivated by its status as an open-source, Python-based library. This offers two distinct advantages: First, its open-source nature ensures transparency and accessibility for the research community. Second, being a Python library, it provides flexibility and can be easily extended. This allows for integration with other scientific computing libraries, such as PyTorch, for machine learning tasks.

At its core, Pandapower provides a tool to model power grids and solve power flows. The library features an implementation of the Newton-Raphson power flow solver. The practical utility of the framework is significantly enhanced by its direct integration with the SimBench grids, which are available directly as Pandapower grids, providing standardized, validated, and realistic network topologies for testing and analysis. This ensures the reproducibility of results.

The modeling process itself is streamlined through the library of standardized component models for power system infrastructures. This includes predefined models for common lines, switches, and transformers, which accelerates the grid modeling process.

As numerous modern publications in the field of energy informatics utilize Pandapower [De+22; Dem+25; Lin+24], it proves to be a highly suitable and powerful tool for the scope of this research.

## 2.5 Power Grid Graph Representation

In the present thesis, the distribution grid is modeled as an undirected graph,  $G = (N, E)$ , where the set of nodes  $N = 1, \dots, n$  corresponds to the electrical buses, and the set of edges  $E \subseteq N \times N$  represents the power lines. Each node  $n$  embodies a distribution grid participant, such as a building or DER. Every node is associated with multiple attributes: a voltage magnitude, a voltage angle, a reactive power, and an active power. Each edge  $e \subseteq E$  is characterized by the admittance of the corresponding power line ( $Y = G + jB$ ). The graph's adjacency matrix is closely related to the admittance matrix. The adjacency matrix indicates the weight of an edge between nodes, while the negative off-diagonal elements of the admittance matrix quantify the ease with which current flows between those buses.

### 2.5.1 Graph Neural Networks

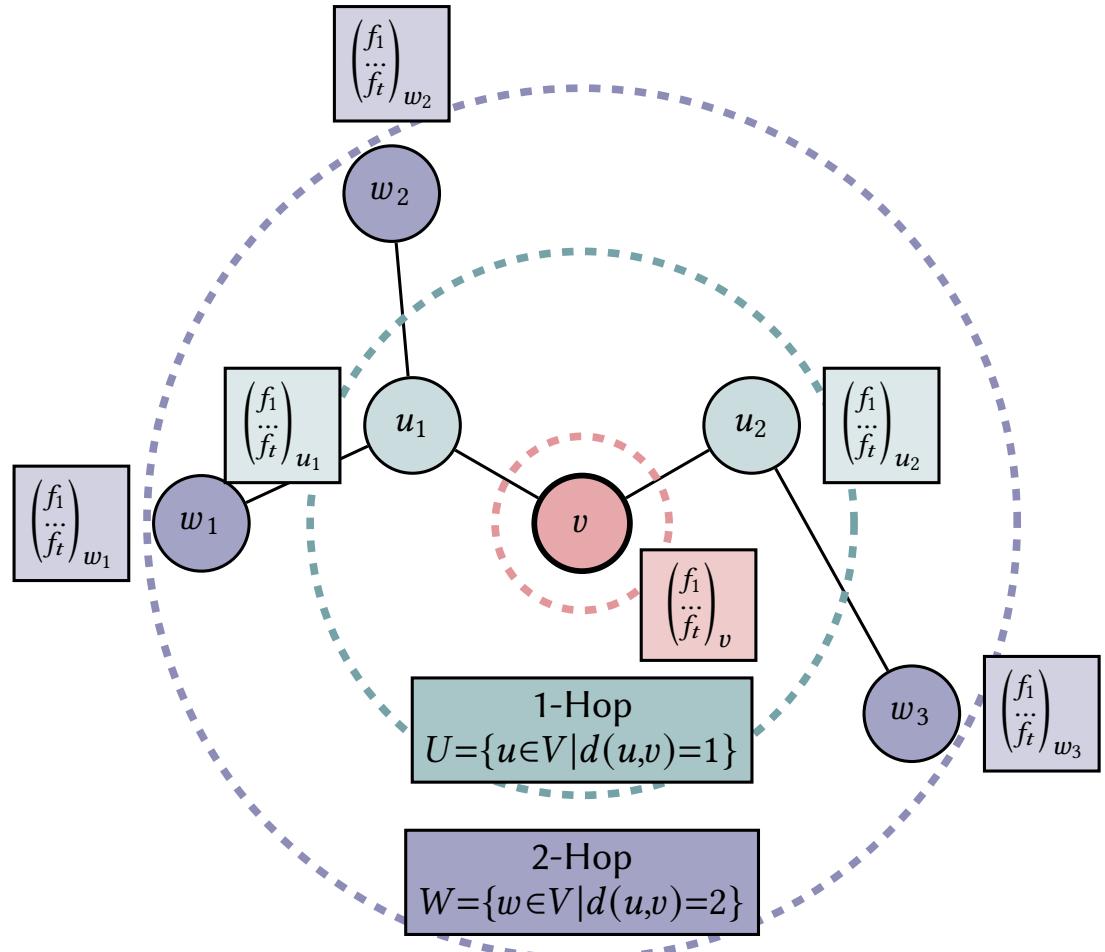
GNNs are a specialized class of deep learning architectures designed to operate directly on graph data. Both the input and the output of these models are based on graphs. A fundamental feature that distinguishes GNNs from traditional architectures such as Convolutional Neural Networks (CNNs) is their permutation invariance. While the performance of traditional models depends on a fixed, ordered arrangement of the input data (e.g., pixels in a grid or words in a sequence), the topological information of a graph is preserved even if the order of its nodes is permuted in the data representation - for example, in the adjacency matrix. A GNN is designed to respect this structural property. Consequently, the model treats different but isomorphic graph representations of the same graph identically and produces a consistent output that depends solely on the graph topology and node features, not on random indexing.

Within the broader landscape of GNNs, Graph Convolutional Networks (GCNs) constitute a prominent and influential subclass. Their central objective is to generalize the concept of CNNs to the irregular, topological structure of graphs. The methodologies for achieving this generalization are primarily divided into two fundamental paradigms: spectral GCNs and spatial GCNs. Spectral GCNs are grounded in the principles of graph signal processing, defining the convolution in the graph's frequency domain. This approach leverages the graph's Laplacian matrix, defined as  $L = D - A$ , where  $A$  is the adjacency matrix of the graph, and  $D$  denotes the diagonal degree matrix. The diagonal elements  $D_{i,i}$  contain the degree of node  $i$ , which is defined as the number of edges connected to node  $i$ . In contrast, spatial GCNs define convolution directly on the graph's topology, operating in the spatial domain. This methodology is directly analogous to the kernels used in traditional GCNs. For each layer, a node's feature representation is updated by aggregating the feature representations of its neighbors. This paradigm can be conceptualized as a message-passing or neighborhood aggregation scheme. For each node, a convolutional operation consists of multiple steps. First, the feature vectors from its  $K$ -hop neighbors are collected. The set of  $K$ -hop neighbors for a vertex  $v$  in a graph  $G = (V, E)$  is formally given by  $n_K(v) = \{u \in V | d(u, v) = K\}$ ,

where the function  $d(u, v)$  denotes the length of the shortest path connecting vertices  $u$  and  $v$ .

The feature vectors are combined into a single vector, typically through a differentiable function like a (weighted) mean or sum. The combination of these feature vectors can also be accomplished with a neural network; for example, the PowerFlowNet model by Lin et al. [Lin+24] employs an Multilayer Perceptron (MLP) for this purpose. Finally, the aggregated neighborhood information is combined with the central node's own current representation.

For illustrative purposes, Figure 2.2 provides a schematic of a simple graph structure. The concept of a node's receptive field is demonstrated for a central node,  $v$ , where its 0-hop, 1-hop, and 2-hop neighborhoods are denoted by colored circles. The feature vector of each node  $\begin{pmatrix} f_1 \\ \vdots \\ f_t \end{pmatrix}$  is depicted next to the respective node. The combination of the neighboring feature

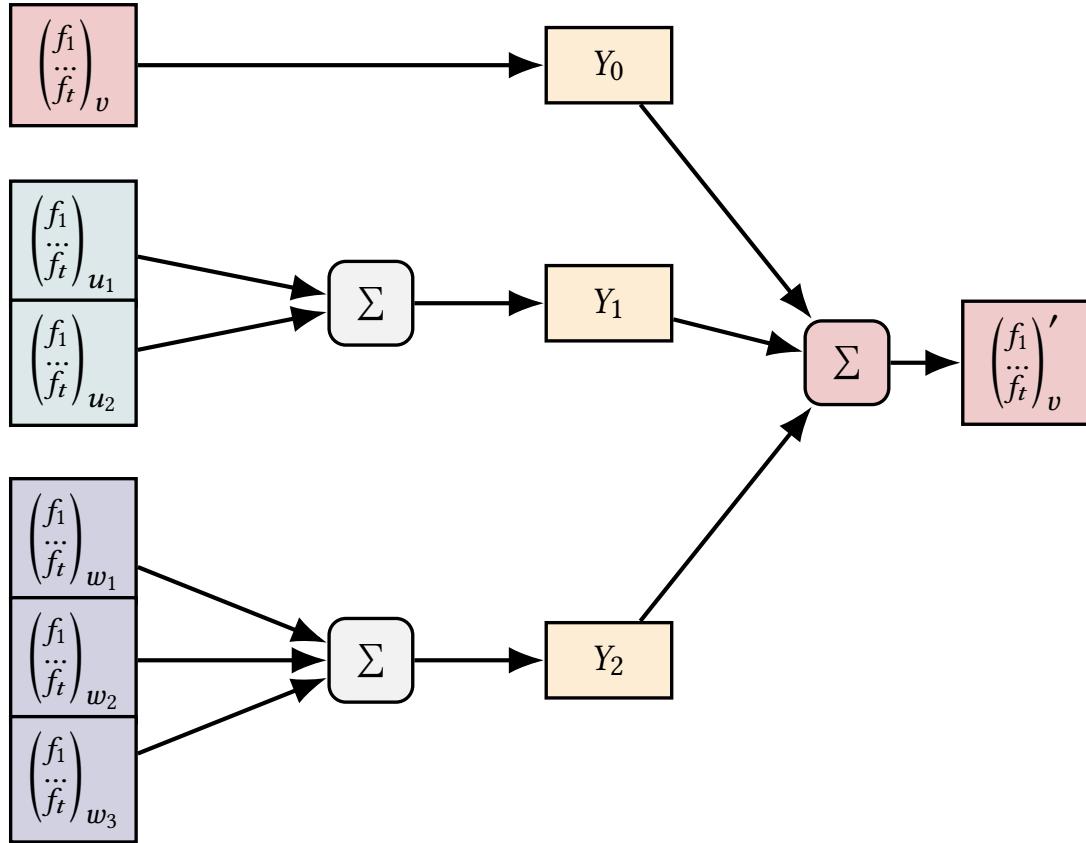


**Figure 2.2:** Exemplary graph with 0-hop, 1-hop, and 2-hop neighborhood of node  $v$ .

vectors and the node's own representation can be achieved using various convolutional operators. The Topology Adaptive Graph Convolutional (TAGConv) operator, introduced in the work of Du et al. [Du+17], produces an output tensor  $\mathbf{X}'$ , represented as

$$\mathbf{X}' = \sum_{k=0}^K \left( \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2} \right)^K \mathbf{X} \mathbf{W}_K. \quad (2.18)$$

In this formulation, the term  $\mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$  represents the symmetrically normalized adjacency matrix,  $\mathbf{A}$ . This is computed using the adjacency matrix  $\hat{\mathbf{A}}$  and its corresponding diagonal degree matrix  $\mathbf{D}$ , where each diagonal element is defined as  $D_{ii} = \sum_{j=0} A_{ij}$ . Furthermore, the parameter  $K$  specifies the number of hops considered within the convolutional layer, and  $\mathbf{W}_K$  denotes the trainable weight matrix that is applied to the input feature tensor  $\mathbf{X}$ . An exemplary convolution with  $K = 2$  is depicted in Figure 2.3.



**Figure 2.3:** Illustration of a graph convolution. The feature vectors are aggregated and incorporated into the graph convolution, depending on the distance.

## 2.5.2 Multilayer Perceptrons

MLPs consist of multiple neural network layers with activation functions and no feedback connections. They are also referred to as feedforward neural networks and approximate a function  $y = f(x; \theta)$  with learnable parameters  $\theta$ . Since each layer represents a function, a three-layer multilayer perceptron with the layers being connected in a chain can be expressed as  $y = f^3(f^2(f^1(x)))$  [GBC16].

## 2.5.3 Training Procedure

The gradient descent method is an iterative optimization algorithm used to find a local minimum of a differentiable function. Within the domain of machine learning, the objective is to methodically adjust the parameters of a model, such as the weights and biases of a neural network, in a manner that minimizes a specific loss function.

This function is employed to quantify the error or “loss” of the model for a specified set of parameters. It measures the discrepancy between the model’s predictions and the actual target values. In the present thesis, the MSE is utilized as either the loss function  $J$  or as a component of the loss function. The objective is to find the parameter values that minimize the loss function. The adjustment of the model parameters depends on the learning rate  $\alpha$ . This scalar hyperparameter determines the length of the step taken during each iteration, corresponding to the amount of progress towards the optimal parameters. A small learning rate results in slow but stable convergence, while a large learning rate can speed up the process but risks overshooting the minimum [Jun22]. In order to adjust the parameters in the correct direction, the back-propagation algorithm computes the gradient. This approach enables the information from the loss function to flow backwards through a straightforward procedure [GBC16].

The weights and biases are updated based on the gradient. For the gradient descent method, the weights  $w$  are updated according to

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \alpha \nabla f(\mathbf{w}_t). \quad (2.19)$$

The term  $\nabla f(\mathbf{w}_t)$  in Equation 2.19 represents the gradient of the loss function that is multiplied by the learning rate  $\alpha$  to modify the weights  $\mathbf{w}$ .

In the present thesis, the Adam optimizer is employed. The name “Adam” stems from the adaptive moment estimation. The algorithm accelerates the gradient descent process by integrating a moving average of the first- and second-order moments  $\mathbf{m}$  and  $\mathbf{v}$  of the gradient [KB17]. The first-order moments  $\mathbf{m}$  are updated according to Equation 2.20. The parameters  $\beta_1$  and  $\beta_2$  represent the exponential decay rates that quantify the influence of the gradient on the moments. They define the extent to which preceding moments influence subsequent ones. According to the original publication of the Adam optimizer, “good default settings” are  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$ . These settings are therefore used in the present thesis. Note that  $\beta_1 = 0$  would result in the gradient descent described above (Equation 2.19):

$$\mathbf{m}_{t+1} = \beta_1 \cdot \mathbf{m}_t + (1 - \beta_1) \cdot \nabla f(\mathbf{w}_t). \quad (2.20)$$

The second-order moments  $\mathbf{v}$  are updated according to

$$\mathbf{v}_{t+1} = \beta_2 \cdot \mathbf{v}_t + (1 - \beta_2) \cdot (\nabla f(\mathbf{w}_t))^2. \quad (2.21)$$

Both moments  $\mathbf{m}$  and  $\mathbf{v}$  are bias-corrected using the terms  $\hat{\mathbf{m}}_t = \mathbf{m}_t / (1 - \beta_1^t)$  and  $\hat{\mathbf{v}}_t = \mathbf{v}_t / (1 - \beta_2^t)$ , respectively, as they start out at zero and tend to be biased towards zero. The weights are finally updated using

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \alpha \frac{\hat{\mathbf{m}}_t}{\sqrt{\hat{\mathbf{v}}_t + \epsilon}}. \quad (2.22)$$

In the Equation 2.22 to update the weights,  $\epsilon = 10^{-8}$  is another hyperparameter that is required to avoid division by zero.



The operation of modern distribution grids is undergoing a fundamental transformation, driven by the increasing penetration of renewable generation at low-voltage levels, the proliferation of EV charging infrastructure, and the widespread adoption of heat pumps for residential heating. While critical for the energy transition, these new assets introduce unprecedented levels of stress and volatility into grids that were traditionally designed for unidirectional power flow. Consequently, there is a growing risk of violating established operational guidelines and standards. To maintain grid stability and reliability, a transition from reactive to proactive grid operations is essential. This requires the development of new monitoring techniques capable of providing real-time situational awareness and detecting potential violations before they occur. This chapter investigates the viability of using measurement data from smart home hardware as a cost-effective solution to this challenge, motivating the first research question: **To what extent can smart plugs be used for distribution grid monitoring?**

After collecting and pre-processing the raw measurement data from the smart plugs, the generation of pseudo-measurements is essential for achieving complete distribution grid monitoring. The evaluation of the accuracy of these generated pseudo-measurements addresses the second research question: **How viable is the use of a feature propagation algorithm for pseudo-measurement generation in distribution grids compared to other related methods?**

This chapter is primarily based on the publication “Distribution Grid Monitoring Based on Widely Available Smart Plugs” [GFH24b] and “Distribution grid monitoring based on feature propagation using smart plugs” [GFH24a].

### 3.1 Introduction

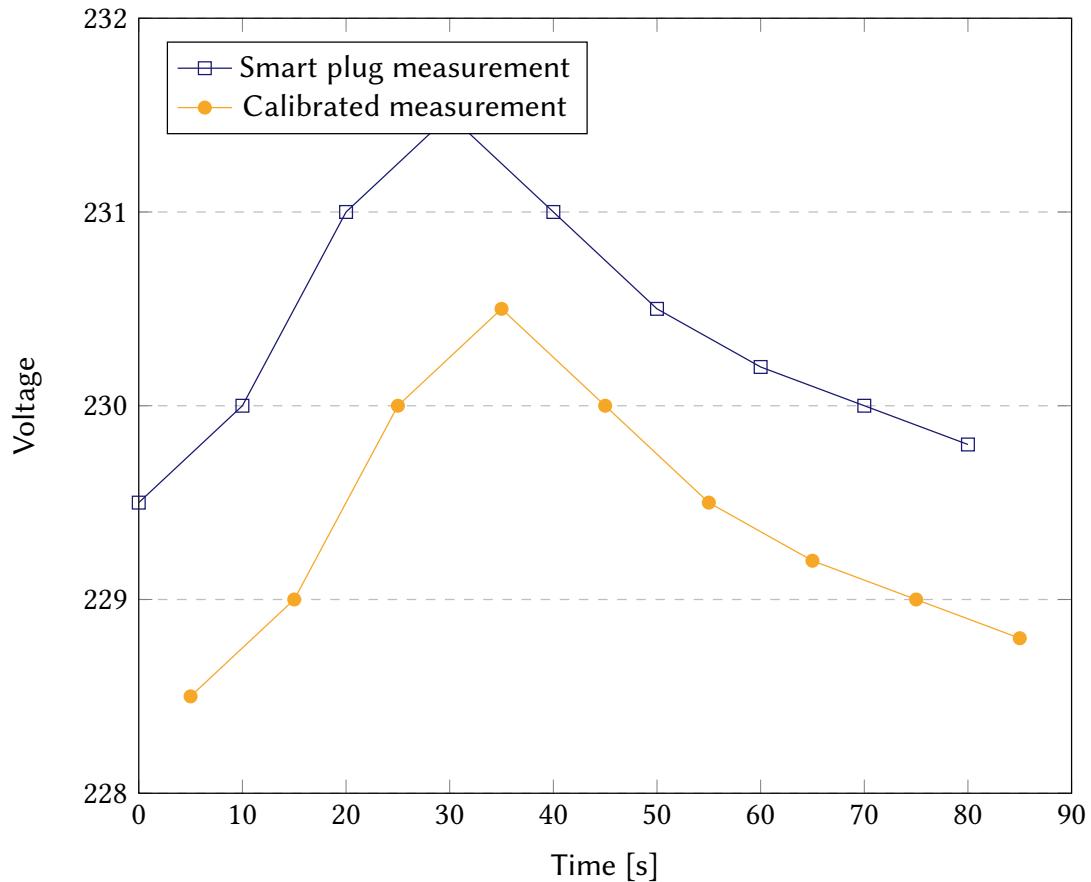
The voltage in European distribution grids must be maintained within the range of 0.9 to 1.1 p.u., as defined in the EN-50160 standard. Voltage levels are influenced by line impedances, with Ohm’s law dictating that voltage drops across lines are proportional to current strength ( $U = Z \cdot I$ ). Consequently, higher currents lead to greater voltage drops. The integration of high-power loads, such as EV charging stations and heat pumps, increases the loads and thereby further reduces the voltage levels. On the other hand, voltage increases can occur due to decentralized generation systems, such as PV installations. The distribution grid has a finite capacity to accommodate decentralized generation without exceeding voltage

limits, a concept referred to as hosting capacity [FPL20; MBE20].

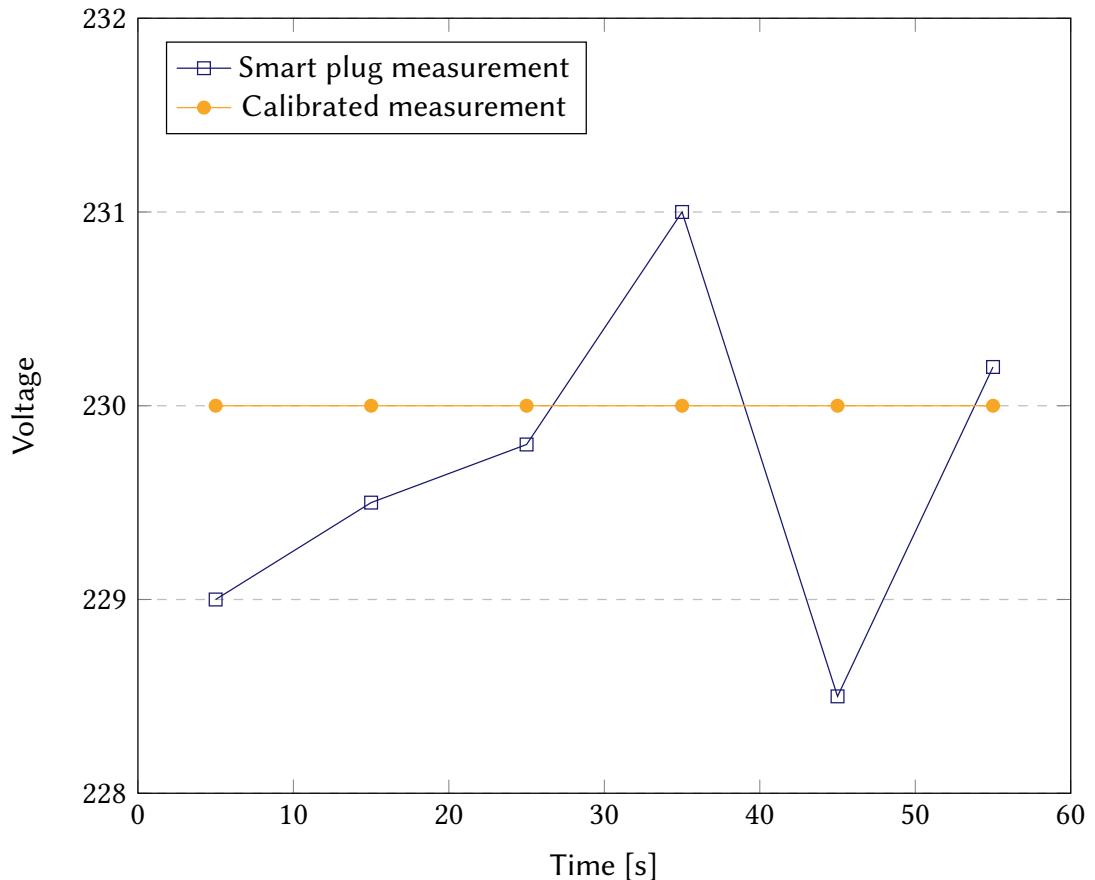
However, the new loads and generators do not operate at constant power levels. They are heavily influenced by external factors such as weather conditions and user behavior, which can cause them to increase or decrease their input and output rapidly. As a result, the power grid experiences rapid voltage fluctuations, which can have detrimental effects on the overall stability and reliability of the system. DERs, such as controllable heat pumps, schedulable EV chargers, and DSM of smart household appliances, have the potential to mitigate these voltage fluctuations when they are detected. However, effective mitigation requires monitoring and responding to changes in the grid, necessitating fast and high-frequency measurements. The DSO needs to have comprehensive and up-to-date measurement data of the whole distribution grid to control these systems. Moreover, the design and optimization of new DERs rely heavily on the analysis of historical voltage data, which provides valuable insights into the patterns and trends of voltage fluctuations in the grid.

The acquisition of accurate and timely measurements is crucial for effective distribution grid management. However, the deployment of smart distribution grid infrastructure and measurement devices is proceeding slowly due to the prevalence of aging network infrastructures. These legacy systems lack monitoring and communication capabilities, and upgrading them is a complex and costly endeavor. The lengthy lifespan of critical infrastructure components, including transformers and cables, which is typically in the range of 30 years [ASB11; Bar+95], introduces an additional layer of complexity to the modernization process. Premature replacement of these components before the end of their operational lifespan is often economically unfeasible. Consequently, this underscores the necessity for novel, cost-effective solutions and strategic investment approaches to facilitate grid modernization, thereby mitigating the financial burdens associated with infrastructure upgrades.

Currently, the adoption of smart household devices is experiencing rapid growth. According to recent studies, approximately 35.8 % of households in the United States and 23.0 % of households in the European Union had smart home devices installed as of 2021 [Mar22], with a projected annual growth rate of around 9 % [Mar24]. A specific category of smart home devices, known as smart plugs, is designed to be interposed between an appliance and a socket outlet, enabling remote control of the connected device's power supply. Some smart plugs can measure the power consumption of connected devices through voltage and current measurements. Modifying the firmware of these smart plugs enables reconfiguration to communicate with a custom server, rather than the manufacturer's default server, and transmit measured data to this alternative server. It appears that smart plugs have the potential to provide voltage measurements with sufficient accuracy for distribution grid monitoring applications. This suggests a possible opportunity for leveraging these devices in the context of grid monitoring.



**Figure 3.1:** Exemplary depiction of the constant offset of the smart plug voltage magnitude measurements. Originally published in [GFH24b]



**Figure 3.2:** Exemplary depiction of the fluctuating error of the smart plug voltage magnitude measurements. Originally published in [GFH24b]

The raw voltage data acquired from the smart plug instrumentation is subject to measurement error. This error can be broken down into two components: a systematic bias and a stochastic error. The systematic bias manifests as a constant offset, introducing a fixed deviation to all measurements, independent of the true voltage magnitude. This type of error is caused by persistent effects, such as offsets from factory calibration or the tolerance of passive components like voltage-divider resistors. A conceptual representation of this constant offset is provided in Figure 3.1. The second error component is a varying error that introduces non-systematic fluctuations in the measured data. This random error can include thermal noise within the electronic components, quantization errors due to the finite resolution of the analog-to-digital converter, and potential artifacts introduced by signal interpolation. This varying error is illustrated in Figure 3.2. To improve the fidelity of the measurements, a calibration procedure can be performed to mitigate the systematic bias. By monitoring the smart plug measurements for a prolonged duration, the constant offset can be quantified. A correction can then be applied by subtracting this calculated offset from the measured voltage magnitude. The resulting corrected measurements are

effectively centered on the true value, with the residual uncertainty being characterized by the varying stochastic error.

Yet, even using widely available smart plugs as measurement devices, complete grid coverage is hard to achieve. Therefore, it is necessary to substitute missing measurement data, for example, with pseudo-measurements. These artificial measurements can be derived from historical measurement data or from data collected by existing measurement devices. However, several requirements must be fulfilled to ensure the practicability of pseudo-measurements. The generation of these measurements must be rapid, enabling a timely response to fast fluctuations in the grid. This implies that the computational effort required for generating the measurements needs to be relatively low. Furthermore, the produced measurements must exhibit a level of accuracy that is sufficient for power grid monitoring purposes. For established techniques, historical data or typical load profiles are often necessary [Ang+16; Man+12; SPJ10; Zha+23]. Due to the ever-changing distribution grid infrastructure and connected appliances, relying solely on historical data may not yield accurate results for the current grid state. Therefore, deriving pseudo-measurement data from the available measurements without the inclusion of historical data is advantageous. Graph-based algorithms can be used for such tasks, as power grids can be represented as graphs. In the test cases considered in the present work, graphs representing distribution grids are cycle-free. Furthermore, the graph's nodes representing the transformers, households, and other connected facilities in the grid have homophilic characteristics. These features are the voltage and voltage angle, which are very similar at neighboring nodes. Thus, algorithms exploiting these homophilic properties can be applied.

The heat diffusion equation models the heat transfer within a material due to a temperature gradient. Fourier's law of conduction describes this relationship with

$$Q \propto A \frac{\partial T}{\partial x}. \quad (3.1)$$

In Equation 3.1,  $Q$  is the heat transferred,  $A$  is the area over which the heat is transferred, and  $\frac{\partial T}{\partial x}$  is the temperature gradient [Geo24].

When discretized, the material exhibits homophily, with neighboring elements having similar temperatures. Rossi et al. have developed a graph feature propagation algorithm based on the heat diffusion equation, demonstrating its effectiveness in generating accurate features in homophilic graphs through neighbor-based aggregation [Ros+21].

In the present thesis, smart plugs provide the available measurement data. They measure the voltage at a single phase at a standard power outlet, and the feature propagation algorithm is used to generate voltage pseudo-measurements at nodes that are not instrumented. However, without loss of generality, this method can be applied to all three phases of the grid, enabling comprehensive monitoring.

This chapter commences with a comprehensive review of the relevant literature, first covering state-of-the-art measurement devices in distribution grids. The review then delves into methods for enhancing the state estimation, generating pseudo-measurements, and imputing missing data. Following this literature review, the chapter details the proposed methodology, beginning with the data collection process, followed by an explanation of the feature propagation algorithm. The subsequent evaluation is twofold: first, the intrinsic accuracy of the raw smart plug measurements is assessed. Second, leveraging the measured data, the efficacy of the feature propagation approach for generating robust pseudo-measurements is evaluated with varying amounts of missing values. The chapter concludes by synthesizing these findings and exploring their broader implications for the field of distribution grid monitoring.

## 3.2 Related Work

Accurate voltage measurements at multiple nodes in the grid are the foundation for assessing the state of the distribution grid [And+11; FPL20]. For example, calculations of the hosting capacity are based on the voltage levels within the grid. The hosting capacity determines the extent to which renewable generation systems such as solar panels and wind turbines can be integrated into the existing power grid without threatening its stability, reliability, or power quality [EB11]. Because the hosting capacity of PV systems depends on many factors, such as grid layout, topology, connection type, and others, individual measurements do not allow conclusions to be drawn about an entire grid area. When assessing the hosting capacity in a distribution grid area, it is essential to determine the minimum and maximum voltage levels that occur over a specified time frame, which necessitates conducting long-term voltage measurements.

Previously, the absence of measurement hardware in the distribution grid led to simulations relying on limited measurement data and pseudo-measurements [AAM22; Dah+20; PL17]. However, the integration of smart meter measurement data into the state estimation can also increase observability in the distribution grid [Ali+14; JCL13; KH22; Wan+19].

The measurement fidelity of smart meters has been demonstrated to be sufficient for enhancing load forecasting accuracy [Qui+15], improving state estimation precision [Sam+11], and detecting errors and faults within the grid infrastructure [Ara+19; TF17]. To address the challenge of grid state estimation under conditions of limited observability from a limited number of smart meters, network topology reduction can be employed [KH22]. Furthermore, the inherent asynchronicity of measurements from multiple smart meters presents a challenge for accurate state estimation. To mitigate this issue, a preprocessing methodology can be employed wherein individual measurements are weighted prior to estimation [Ali+14]. In particular, the weights assigned to each measurement are calculated as a function of the measurement's age. This technique has been demonstrated to yield superior accuracy compared to traditional state estimation algorithms, which erroneously assume simultaneous data acquisition. Nonetheless, professional electricians are needed to install smart meters, and these meters typically capture measurements at 15-minute intervals [Wan+19].

Utilizing this already existing Advanced Metering Infrastructure (AMI) in the distribution grid results in cost savings and reduced expenditures, albeit at the expense of data timeliness and, consequently, the immediate accuracy of grid state estimation [San+17]. The absence of measurement hardware also results in imprecise distribution grid transformer loading data. AMI can be incorporated into analyses to calculate transformer load profiles and assess whether the current infrastructure could support potential new loads without risk of overloading [LHL14]. While installing monitoring devices on all transformers could address this issue, it is not economically feasible to replace working infrastructure [LHL14].

Additional information sources, such as temperature, geographic, customer, and facility management data, can be integrated to enhance grid models and more accurately estimate transformer peak loads. The near real-time optimization of the distribution network through smart grid technology represents a substantial advancement for the efficient operation of the grid [LHL14].

In the research conducted by Dimitriou et al. [Dim+14], a smart plug capable of real-time voltage and frequency monitoring is developed. The measured data is transmitted via Bluetooth to a smartphone, which subsequently uploads the information to a web server. This implementation illustrates the feasibility of measuring voltages at various household outlets within the distribution grid and estimating the grid's real-time status based on these measurements. The device serves as a proof of concept for a cost-effective alternative to traditional AMI hardware. However, the study does not specify the measurement accuracy or any associated time delays. In their 2011 study, Abart et al. put forth a proposal for the utilization of specialized voltage meters for the purpose of monitoring the electrical grid [And+11]. Their proposed methodology incorporates a synchronization of measurements and the execution of grid state analyses through the implementation of load flow simulations based on sequential snapshots of the grid. Furthermore, the foundational grid model refinement is achieved by incorporating insights derived from discrepancies between the calculated and observed voltages at various nodes [And+11]. In the study by Ganu et al. [Gan+12], a smart plug is designed for DSM. The authors developed software capable of toggling the connected load based on the detected voltage levels, demonstrating that widespread deployment of these smart plugs within the grid allows for effective peak shaving. However, the device lacks a communication mechanism, rendering the measured data unsuitable for distribution grid monitoring. The hardware remains in the prototype stage and is unavailable for commercial purchase. Melo et al. [Mel+21] develop a cost-effective data logger equipped with LoRa wireless communication to monitor meteorological variables and PV generation. LoRa is a modulation technique for the transmission of small data packets using low power, while achieving a long range [bad23; LoR]. This device transmits data to a LoRa Gateway, which then forwards it to an MQTT (a message queue with publishers and subscribers) broker [And+19], with subsequent storage in the Google Cloud Platform. Despite the innovative design, these devices are custom-built and not readily available on a commercial scale, thus limiting their potential for widespread adoption.

Vanting et al. [VMJ21] demonstrate that by integrating various input variables, including historical data, weather conditions, and daily features, the loads on the distribution grid can be accurately forecasted using deep neural networks. These forecasts can then be used to infer the grid's state [VMJ21]. However, due to averaging effects, the aggregation of loads from multiple distribution grid participants results in a less volatile total load compared to that of a single household [VMJ21]. Nevertheless, when attempting to predict the load of a single household, additional external factors must be considered, such as occupancy behavior and building characteristics. Consequently, effectively monitoring a single branch

of the distribution grid with only a few participants necessitates detailed information about the occupants. In the work by Manandhar et al., Kalman filters are employed to estimate various state variables within the distribution grid [Man+14]. Discrepancies between the Kalman filter estimates, derived from a mathematical model of the power grid, and actual measurement data serve as indicators for triggering alarms. In their study, attacks and faults in the grid are detected by comparing predicted and measured states. These predictors provide a thorough estimate of distribution grid measurement values, offering an alternative to traditional monitoring methods [Man+14].

According to Markovic et al. [MBH23], effective monitoring of the distribution grid is contingent upon the availability of sufficient measurement data. For nodes lacking direct measurement hardware, measurement data can be inferred from alternative sources. For instance, the authors employ estimates from preceding time steps to create pseudo-measurements, thereby enhancing grid monitoring capabilities [MBH23]. As Manitsas et al. emphasize, when typical load profiles are available for missing measurement values, they can be integrated into the process of generating pseudo-measurements [Man+12]. However, actual load profiles frequently deviate significantly from these typical profiles, necessitating the application of corrective measures [Ang+16]. Alternative methodologies assume that loads within the distribution grid adhere to probabilistic density functions, facilitating the generation of pseudo-measurements based on these distributions [SPJ10]. Alternatively, multiple data sources can be combined within a machine learning framework to produce accurate voltage pseudo-measurements [Zha+23]. This approach integrates historical data with current measurements, achieving high voltage estimation accuracy and performing computations in a matter of milliseconds [Zha+23]. The method demonstrates particular robustness in scenarios characterized by high penetration of PV systems.

A feature propagation algorithm is proposed in [Ros+21] that enables the generation of missing data features without reliance on prior time steps, conventional load profiles, or statistical load distributions. This approach utilizes a graph structure to diffuse known features, thereby reconstructing the missing features. Rossi et al. [Ros+21] conceptualize feature propagation as a low-pass filter, stating that it is particularly well-suited for homophilic graph data. In homophilic graphs, adjacent nodes exhibit similar features, a characteristic that applies to voltages in the distribution grid, where the voltage values of neighboring nodes tend to be closely aligned [Cos+11]. To evaluate the effectiveness of feature propagation, node classification tasks are conducted using a GCN on various benchmark datasets. A GCN is a type of GNN that exhibits invariance or equivariance with the input data. Consequently, permutations of the input nodes either leave the output nodes unchanged or permute them correspondingly. This characteristic enables a GCN to generalize across all node permutations within a graph [KP19]. Incorporating both the propagated and original features into a GCN, the method proposed by Rossi et al. results in only a marginal accuracy reduction of 4 % even with 99 % of features missing [Ros+21]. They demonstrate in their evaluation that this feature propagation approach surpasses four

other methods for reconstructing missing features on benchmark datasets, including Cora, Citeseer, and PubMed.

Spinelli et al. propose an alternative methodology for estimating missing values within a dataset [SSU20]. They construct a similarity graph from the training data to represent the connections between individual features, with the edges' weights quantifying the similarity between nodes. This approach allows for the creation of a graph representation of tabular data. The resulting similarity graph is then encoded into a higher-dimensional intermediate graph using a GCN autoencoder. Subsequently, another GCN decodes this intermediate representation back into the original data dimension, incorporating the imputed features. In comparison to other data imputation algorithms, this method achieves superior results in most instances.

In the present thesis, an approach for monitoring distribution grids through feature propagation is presented, first outlined in [GFH24a]. The feature propagation algorithm is heavily influenced by the work of Rossi et al. [Ros+21]. Unlike machine-learning techniques [Man+12] or probabilistic mixture models [SPJ10] that are traditionally employed to generate pseudo-measurements, the presented method leverages the homophilic properties of nodes within the distribution grid. This approach obviates the necessity of defining typical load profiles for grid participants and mitigates concerns about substantial deviations from these profiles, which can occur due to the increasing integration of renewable generation components. To our knowledge, this feature propagation method has not previously been applied to electricity grids in the existing literature. The Graph Imputation Neural Network (GINN) algorithm presented by Spinelli et al. [SSU20] serves as a benchmark for comparison with the approach presented in the present thesis.

### 3.3 Method

This section details the methodology designed to address the first research question: **To what extent can smart plugs be used for distribution grid monitoring?** To this end, a monitoring framework utilizing smart plug measurements is implemented. However, practical constraints prevent the installation of smart plugs at every node within the distribution grid, resulting in sparse measurement data. To overcome this limitation and ensure network observability, a pseudo-measurement generation technique is employed to augment the real data. To evaluate the performance of this technique, it is benchmarked against a state-of-the-art GINN algorithm. This comparative analysis directly addresses the second research question: **How viable is the use of a feature propagation algorithm for pseudo-measurement generation in distribution grids compared to other related methods?**

#### 3.3.1 Measurement Hardware

Smart plugs contain various measurement hardware. Commonly employed energy measurement integrated circuits (ICs) include the Shanghai Belling BL0937 and BL0940 [Tas22]. The Nous A1T, Gosund SP1, and Shelly Plug S smart plugs utilized in this study contain the BL0937 Integrated Circuit (IC). The prevalence of this IC across a multitude of smart plug models suggests that the measurement accuracy of these smart plugs is uniform, and consequently, the tested smart plugs exhibit consistent measurement characteristics.

Variations in calibration across the smart plugs are observable. Manufacturing tolerances and disparate environmental conditions experienced by individual units contribute to inter-device variability in measured voltage and current levels, even within identical product models and production batches. To mitigate this, a constant offset bias is calculated for each smart plug independently. This bias is determined by quantifying the mean deviation between voltage measurements obtained from the smart plug and those acquired from a calibrated reference meter, a Janitza UMG 604EP-PRO power analyzer [Jan19]. Subsequently, the determined offset voltage is subtracted from the raw voltage measurements during a pre-processing stage. Following this correction, the mean voltage values obtained from all smart plugs subjected to a common, stable voltage source exhibit convergence. Before operational deployment, the initial offset voltage calibration must be conducted for each individual smart plug.

#### 3.3.2 Measurement Data Acquisition

Smart plugs can interface with various network technologies, including Zigbee, a low-power personal area wireless network, LoRaWAN, and Wi-Fi. Each of these technologies exhibits distinct performance attributes, including strengths and limitations, which are summarized in Table 3.1. The smart plugs can be actuated remotely through smartphone applications or

a smart home hub, and some smart plugs contain hardware for measuring the electrical power consumption of connected devices, as well as the line voltage of the socket outlet. The resulting measurement data is generally transmitted to a server maintained by the device manufacturer. This infrastructure enables users to monitor the measured parameters through a web service.

Protocol	Hub needed	Data Rate	Range	Prerequisite
Wi-Fi	no	high	low	numerous access points
LoRaWAN	yes	low	high	few access points
Zigbee	yes	medium	medium	some access points

**Table 3.1:** The selection of appropriate communication protocols for smart home devices is determined by factors including grid topology and internet connectivity. Source: [Son+17].

However, with customized firmware, certain smart home devices have the capability to connect to IoT gateways other than the manufacturer's server. These gateways can transmit the collected data, encapsulated in standardized message formats, to a message broker. This enables the remote monitoring and logging of voltage levels and power consumption for connected devices. The primary objective of smart home devices is not grid monitoring, and manufacturers typically omit specifications regarding the accuracy of their measurement data. Moreover, the potential for firmware modifications introduces a variable that can affect both the precision and sampling frequency of the measurements.

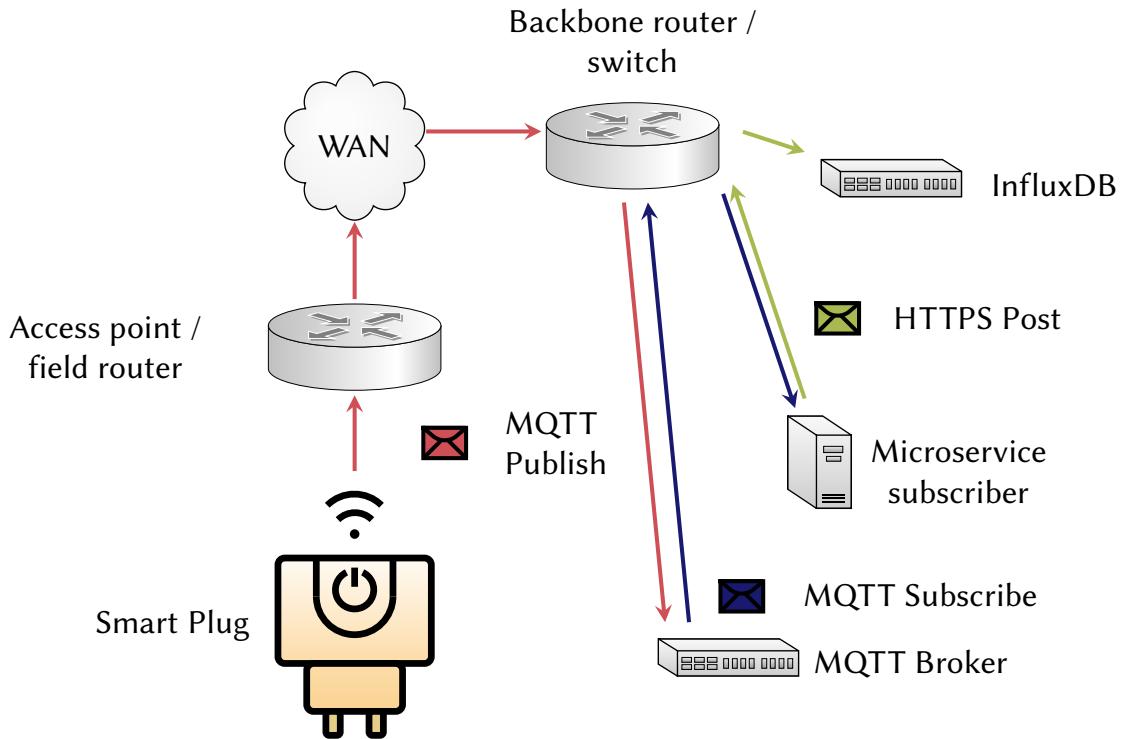
This study employs Wi-Fi-enabled smart plugs reprogrammed with a tailored implementation of the Tasmota open-source firmware<sup>1</sup>. This modification enables the devices to capture voltage measurements at a one-second sampling rate. The acquired voltage data is subsequently transmitted to an MQTT broker.

The selection of Wi-Fi-enabled smart devices, as opposed to LoRaWAN or Zigbee, obviates the requirement for an additional gateway, leveraging pre-existing Wi-Fi access points. This presents a pragmatic deployment strategy within the current experimental framework and anticipated real-world applications, characterized by omnipresent Wi-Fi access points. Besides, the inherent higher bandwidth of Wi-Fi facilitates the transmission of high-frequency, comprehensive measurement data streams. Specifically, the smart plug utilized in this study exhibits a data latency typically below 100 ms, enabling near real-time data accessibility within the simulated environment.

The transmission of data via Wi-Fi is potentially not encrypted. However, the smart plugs employed in this study, which incorporate the ESP8266 microcontroller, support the WPA2 encryption standard to secure communication between the device and the access point. This,

<sup>1</sup> <https://github.com/si-gr/Tasmota/commit/0d5170928e4fc09a29421b43731631c5c061f42e>

in turn, safeguards the integrity of the measurement data transmitted. Furthermore, the Mosquitto MQTT broker utilized in this research supports Transport Layer Security (TLS) encryption. To secure the data transmission from the smart plugs, a customized Tasmota version containing the BearSSL library is created. This lightweight SSL library enables the ESP8266 to perform TLS-encrypted communication. The configuration of the smart plugs with the Secure Sockets Layer (SSL) fingerprint of the MQTT broker and a pre-shared key serves to mitigate the risk of spoofing attacks, rendering the injection of malicious or fabricated data significantly more challenging.



**Figure 3.3:** Network infrastructure between the smart plug and the InfluxDB server. Illustration adapted from [GFH24b].

A microservice is implemented as an adapter between the MQTT broker and an InfluxDB time-series database, with the microservice translating MQTT message payloads into structured data suitable for an InfluxDB storage. This mapping process enables the assignment of measured values to pre-defined fields within the database schema. The architecture supports the concurrent integration of diverse measurement devices, thereby enabling comparative analysis of data, such as voltage readings, across multiple sources. Furthermore, the inclusion of metadata tags within the data stream enables the contextualization of measurements, associating them with attributes such as power phase, geographic location, and device manufacturer. This modular design promotes scalability and extensibility, facilitating the rapid integration of novel sensors through the development of specialized

microservices. The system's flexibility extends to accommodating devices employing alternative communication protocols; new microservices can be developed to bridge these protocols with the time-series database, ensuring backward compatibility. The architecture also permits the horizontal scaling of the data storage by integrating additional database servers, each managed by a dedicated microservice, without requiring modifications of the existing infrastructure. A graphical representation of the network topology is depicted in Figure 3.3.

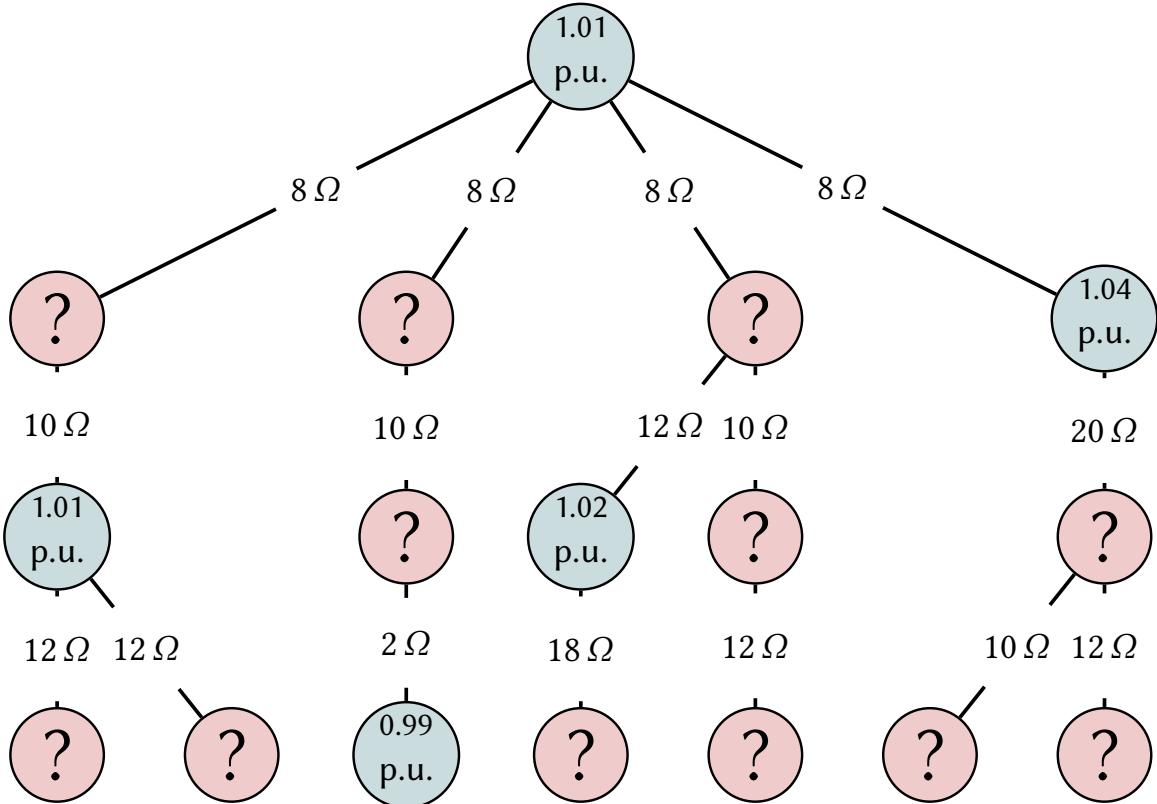
### 3.3.3 Data Processing

While the raw data stream from the smart plugs offers high temporal resolution and low latency, it is insufficient for comprehensive grid monitoring in its raw form. Achieving a complete and observable system state requires a two-stage process. First, the raw measurements undergo a pre-processing pipeline, as detailed in Section 3.3.1. Second, pseudo-measurements must be generated for all uninstrumented nodes to address the inherent data sparsity.

To construct a comprehensive representation of the distribution grid area instrumented with smart plugs, a methodology for generating synthetic measurements is employed. This process utilizes a feature propagation algorithm, adapted from the work of Rossi et al. [Ros+21]. Given the graph-based nature of this algorithm, the initial step involves the construction of a graph representation of the distribution grid topology. The subsequent sections detail the development of the graph data structure and the implementation of the feature propagation algorithm.

### 3.3.4 Feature Propagation

In Figure 3.4, the input data for the feature propagation algorithm is depicted. Nodes colored in red lack associated voltage attributes due to the absence of instrumentation at these locations within the distribution grid. Various strategies exist for imputing these missing features. A rudimentary baseline approach assigns the mean value of the observed voltage values to all uninstrumented nodes. However, this method disregards the inherent topological structure of the network and the underlying data distribution. To address this limitation, a feature propagation algorithm, analogous to heat diffusion, is implemented. This algorithm propagates features based on the adjacency matrix and edge weights. Specifically, feature propagation is achieved through the matrix multiplication of the adjacency matrix with the feature matrix, resulting in an updated feature matrix, effectively propagating features by one hop. The edge weights, representing line conductances, quantify the strength of connectivity between adjacent nodes. During each propagation iteration, the observed voltage attributes are reinstated to their original values, preserving the integrity of empirical measurements and their associated errors. This iterative process, reminiscent of



**Figure 3.4:** Illustration of the distribution grid monitoring with missing features. Nodes colored in blue denote the instrumented nodes, where per-unit measurements are available. Nodes marked in red lack voltage features, necessitating reconstruction. Resistances  $R$  measured in  $\Omega$  represent the edge features that make up the admittance matrix with  $G = \frac{1}{R}$ . The Figure is originally published in [GFH24a].

heat diffusion originating from the instrumented nodes, generates a gradient between these nodes, with the gradient's slope determined by the connection strength. This procedure is detailed in Algorithm 1.

In the feature propagation algorithm,  $x_k$  contains the known features, and  $A$  represents the weighted adjacency matrix. Given the propagation of real-valued voltage measurements and no voltage angles within a modeled DC grid, the adjacency matrix  $A$  is equivalent to the simplified DC admittance matrix. The set  $k = k_1, k_2, \dots, k_n$  represents the indices corresponding to the known features. The vector  $y(i)$  represents the feature vector resulting from iteration  $i$  of the algorithm. Each iterative step necessitates a single matrix multiplication, which can be computed efficiently. The matrix multiplication  $y(i) = A \times y(i - 1)$  is analogous to Fourier's law of conduction as described with Equation 3.1. Just as heat diffuses through a medium based on conductivity and temperature gradients, this algorithm

---

**Algorithm 1** Feature propagation algorithm as presented by Rossi et al. [Ros+21].

---

```

1: procedure FEATUREPROPAGATION( $x, A$ )
2:    $y \leftarrow [0, 0, \dots, 0]$ 
3:    $y_k(0) \leftarrow x_k$                                  $\triangleright$  Set known features
4:   while  $|y(i-1) - y(i)| > \epsilon$  do            $\triangleright$  While not converged
5:      $i \leftarrow i + 1$ 
6:      $y(i) \leftarrow Ay(i-1)$                           $\triangleright$  Propagate features according to adjacency matrix
7:      $y_k(i) \leftarrow x_k$                              $\triangleright$  Reset known features
8:   end while
9:   return  $y(t)$ 
10: end procedure

```

---

propagates voltage features through the grid based on electrical conductance and voltage differences.

The number of iterations required for the convergence of the feature propagation algorithm is influenced by the number and placement of available measurements, as well as the topology of the distribution grid under consideration. The algorithm terminates once the difference  $\delta$  between successive iterations falls below a predefined threshold  $\epsilon$ . For instance, with seven available measurements and a threshold of  $\epsilon = 10^{-5}$ , the algorithm converges after 153 iterations on the IEEE 37-bus system.

In contrast, the GINN algorithm [SSU20], which is used for comparison, does not rely on a graph data structure. Instead, it imputes missing features based on training data, which consists of the same power grid with various loads. While the grid topology remains unchanged within the training dataset, voltage measurements vary according to the applied loads. All voltage measurements are simulated using power flow calculations. The internal GCNs leverage this training data to learn node relationships, eliminating the need for an adjacency matrix as an input.

During each iteration of the feature propagation algorithm, a matrix multiplication operation is executed. Consequently, the computational complexity of this algorithm is inherently dependent on the size of the matrix, which corresponds to the number of nodes in the graph. As a result, the complexity ranges between approximately  $O(n^{2.38})$  and  $O(n^3)$ , depending on the specific algorithm employed for the matrix multiplication [SV23]. In the GINN algorithm, an internal similarity graph is constructed, a process that can be performed efficiently [SSU20]. This graph construction is required only once per dataset during the training phase. However, for the recurrent generation of pseudo-measurements, inference must be performed using GCNs. The complexity of the inference is given by  $O(|E|CF)$ , where  $|E|$  represents the number of edges, and  $C$  and  $F$  denote the number of input and output features of the GCNs, respectively [SSU20].

## 3.4 Evaluation

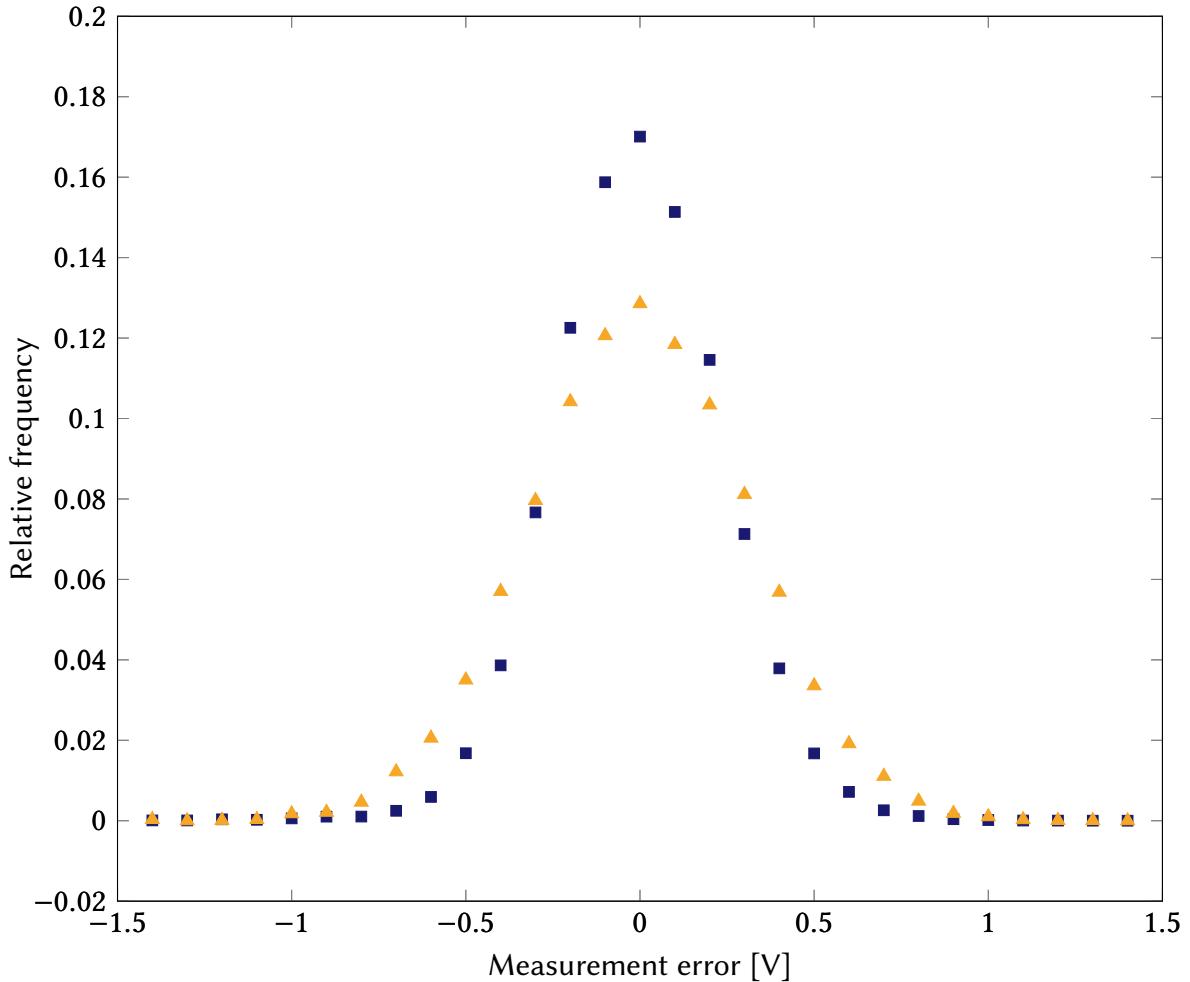
The evaluation of the proposed methodology is structured into two distinct phases. Firstly, the measurement accuracy of the employed smart plugs is assessed. This assessment involves the deployment of two smart plugs within a real-world test environment, where their measured values are compared with those obtained from calibrated, professional-grade instrumentation. Secondly, the effectiveness of the distribution grid monitoring approach is evaluated.

### 3.4.1 Evaluating the Smart Plugs Measurement Accuracy

To assess the accuracy of the voltage measurements obtained from the smart plugs, two devices with identical measurement hardware are deployed within a real-world test environment. The experimental setup involves the configuration of the Janitza power analyzer to transmit measurements at a frequency of one sample per second. Given the utilization of identical measurement hardware within the smart plugs, any discrepancies in the recorded values can be attributed to the variations between the modified and unmodified firmware versions.

Initially, an analysis is conducted on the measurements acquired from the smart plug utilizing the unmodified Tasmota open-source firmware [Tas22]. This firmware version outputs voltage measurements with a precision of one decimal place, suggesting a maximum error of 0.1 V. However, the presence of rounding errors within the unmodified firmware results in a larger measurement error, with the smart plug exhibiting voltage measurement increments of at least 0.2 V, and in some instances, 0.3 V.

Figure 3.5 depicts a comparison of the measurement errors associated with the unmodified and modified firmware versions, colored in orange and blue, respectively. This measurement error is calculated after the preprocessing step and therefore does not contain the constant offset as described in Section 3.3.1. The unmodified firmware transmits measurements at a rate of one sample every ten seconds, whereas the modified firmware provides measurements at a frequency of one sample per second. Consequently, within a given timeframe, the number of measurements obtained from the modified firmware is exactly ten times that of the unmodified firmware. The values represented on the Y-axis are normalized relative to the total number of measurements recorded for each respective firmware version. The measurements acquired from the smart plug utilizing the unmodified Tasmota firmware exhibit greater variability than the distribution observed for the smart plug with the modified firmware. This suggests a lower standard deviation in the measurements associated with the modified firmware (represented in blue) relative to the unmodified firmware (represented in orange). Quantitative analysis verifies this discrepancy, with the standard deviation calculated as 0.33 V for the unmodified firmware and 0.27 V for the modified firmware.



**Figure 3.5:** Relative frequency histogram of the measurement error of smart plugs with the modified firmware (blue) and the unmodified firmware (orange). The depicted measurement error is adjusted to account for the constant offset as described in Section 3.3.1. This Figure was originally published in [GFH24b].

Statistical analysis, employing both the Anderson-Darling test for normality [M A74] and the Shapiro-Wilk test [SW65], failed to reject the null hypothesis of normality for the measurement error data. Specifically, the Anderson-Darling test yielded a test statistic of 0.44, which is less than the critical value of 0.57, even at a high significance level of 15 %. Furthermore, the Shapiro-Wilk test produced a p-value of 0.54. Consequently, based on these results, it can be concluded that the measurement error distribution is consistent with a normal distribution or a closely related distribution.

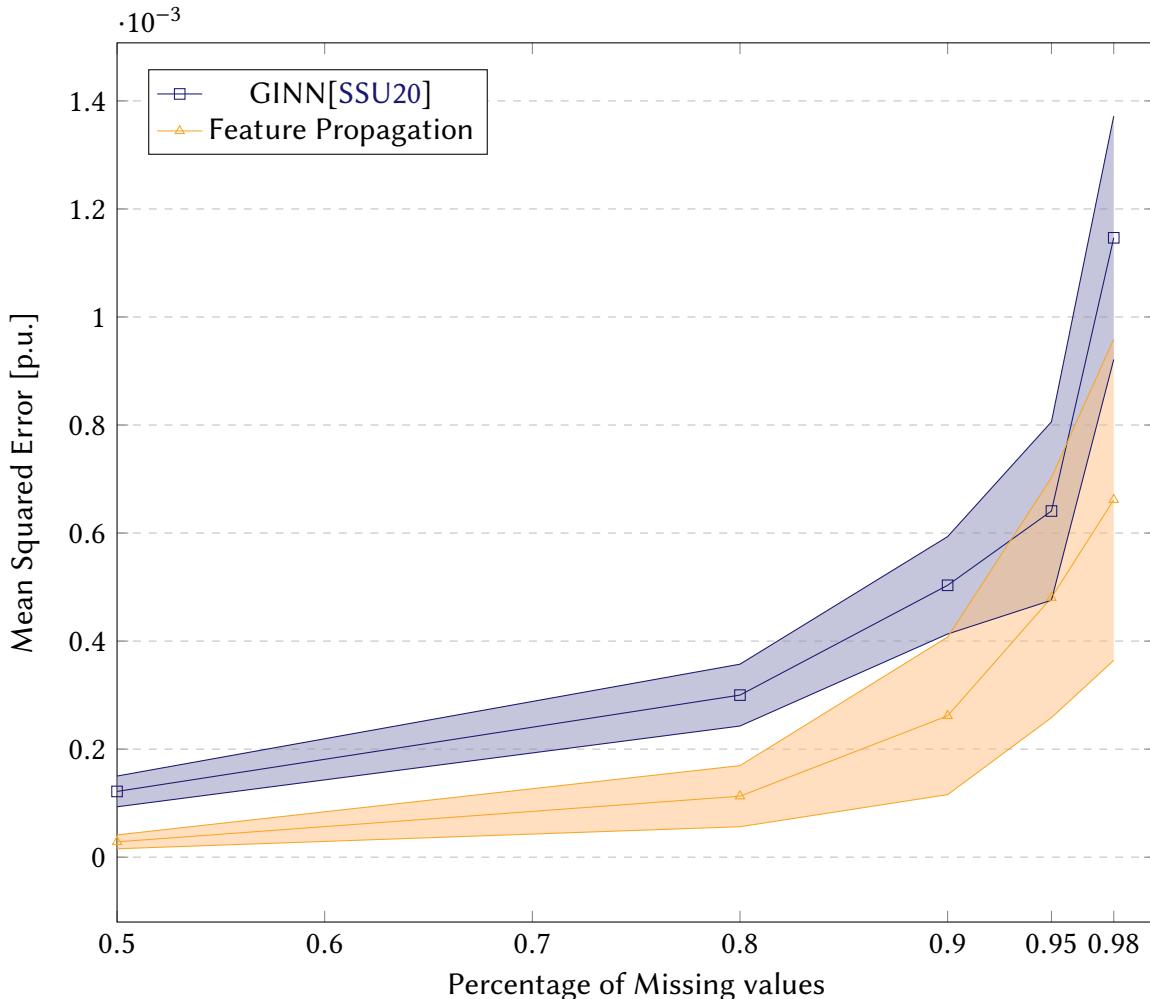
### 3.4.2 Distribution Grid Monitoring Error

For distribution grid monitoring, the Point of Common Coupling (PCC)'s voltage levels are relevant. Regulatory authorities mandate that voltage levels at this demarcation point, between building wiring and the distribution grid operator's network, remain within specified allowable limits. The EN-50160 standard, for instance, defines an acceptable voltage range of 0.9 p.u. to 1.1 p.u. Consequently, the installation of the measurement devices in close proximity to the PCC is essential. This study asserts that the smart plugs utilized for voltage measurement are installed at an electrical outlet located in immediate proximity to the PCC, thereby ensuring that the measured voltage closely approximates the actual PCC voltage. Consequently, the analysis of grid monitoring accuracy can be constrained to the inherent measurement error of the smart plug itself, obviating the need to account for additional errors arising from the device's placement downstream from the PCC.

The present work evaluates the accuracy of distribution grid monitoring through a comparative analysis of simulated monitoring results against known ground truth. The MSE is calculated to quantify the aggregate error. Specifically, the performance of the GINN [SSU20] algorithm was assessed and contrasted with a rudimentary mean imputation technique. This baseline comparison serves to contextualize the GINN algorithm's effectiveness in addressing data gaps within the monitoring framework.

The grid topologies "1-LV-rural13--1", "1-LV-semiurb5--2", and "1-LV-urban6--2" from the SimBench dataset [Mei+20] with corresponding load profiles are used for the simulation. The grids incorporate PV systems at multiple nodes and operate at a nominal voltage of 0.4 kV. They are comprised of 118, 104, and 53 nodes, respectively. Ground truth voltage data is obtained through power flow simulations conducted over 5000 time steps for each grid. To simulate realistic measurement scenarios, a percentage of voltage values was systematically removed, and the remaining values were perturbed with simulated measurement errors drawn from a normal distribution, reflecting the empirical measurement error of the smart plugs. The resulting dataset exhibits varying levels of data sparsity, with 50 %, 20 %, 10 %, 5 %, and 2 % of the voltage values missing (corresponding to 50 %, 80 %, 90 %, 95 %, or 98 % data availability) for each time step and the available voltage values being inaccurate according to the measurement characteristics of the smart plugs.

The performance of the GINN and feature propagation algorithms, in the context of pseudo-measurement voltage generation, is illustrated in Figure 3.6, depicting the MSE with various amounts of missing measurement values. The associated standard deviations are depicted as shaded areas. This evaluation is conducted using 1,000 time steps from the three distribution grid datasets. To simulate realistic measurement scenarios with missing data, a randomized mask is applied to the data at each time step, effectively removing a random selection of input voltage values. The resulting dataset, comprising incomplete voltage measurements for the three grids over 1,000 time steps, is then used as input for



**Figure 3.6:** The MSE between the imputed and true values in the dataset. The shaded area illustrates the standard deviation. Graphic originally published in [GFH24a].

both the GINN and feature propagation algorithms.

To assess the performance of the GINN algorithm, independent training and evaluation procedures are conducted for each of the three distribution grids. This approach, which excludes transfer learning, ensures that the algorithm's performance is evaluated under optimal conditions for each individual grid. Specifically, the GINN algorithm is trained and tested exclusively on data from a single grid, and this process is replicated for each of the three grids included in the evaluation.

Missing Rate	Averaging	GINN	Feature Propagation
0.5	4.375	0.121	0.030
0.8	6.668	0.300	0.105
0.9	8.019	0.503	0.242
0.95	8.792	0.640	0.431
0.98	11.630	1.146	0.626

**Table 3.2:** MSE of the distribution grid monitoring in  $10^{-3}$  p.u. from 1000 time steps of the grids “1-LV-rural3-1”, “1-LV-semiurb5-2”, and “1-LV-urban6-2”. Data originally published in [GFH24a].

The feature propagation algorithm exhibits a high degree of efficacy in reconstructing missing voltage features, maintaining accuracy even under conditions of low numbers of available measurements (95 % and 98 % missing values). While the GINN algorithm’s performance improves with increasing data density, the feature propagation algorithm consistently achieves superior reconstruction accuracy across all tested levels of data sparsity. Quantitatively, with 95 % missing voltage values, the feature propagation algorithm yields an MSE of  $0.43 \cdot 10^{-3}$  p.u., equivalent to 0.099 V in a 230 V grid. Conversely, the GINN algorithm produces an MSE of  $0.64 \cdot 10^{-3}$  p.u., corresponding to 0.147 V in a 230 V grid, under the same sparsity conditions. The naive approach of mean imputation results in substantially higher MSE values. Detailed results are tabulated in Table 3.2.

## 3.5 Discussion

The voltage standard deviation of 0.27 V ( $\approx 0.12\%$ ) achieved by the modified firmware-equipped smart plugs exceeds the performance stated in prior studies, which report acceptable ranges of 0.6 % [Sam+11] and 0.3 % to 0.9 % [JCL13]. The spatial distribution of multiple smart plugs within a distribution grid segment contributes to improved measurement precision. A key advantage of the modified firmware lies in its high temporal resolution, providing voltage measurements at a rate of up to 1 Hz, thereby enabling near-instantaneous grid state evaluation. This contrasts with the comparatively low sampling and transmission rates of traditional smart meters, which often operate with 15-minute measurement intervals and 6-hour data transmission delays [San+17]. The smart plugs also offer significant logistical advantages, with a rapid deployment process that requires minimal technical expertise, enabling uncomplicated field implementation through simple Wi-Fi configuration.

In electrical distribution systems, smart meters are generally installed near the PCC. In contrast, smart plugs are connected to individual outlets and measure the voltage at that point. This difference in installation location leads to a notable difference, as voltage measurements obtained by smart plugs inherently include the voltage drop occurring within the residential electrical grid [GFH23]. Similar to the voltage drop in the distribution grid that is depicted in Figure 2.1, this voltage drop is unique for every household and varies depending on the load conditions within the household's electrical system. It is, therefore, non-constant. To minimize inaccuracies in voltage measurement caused by this internal voltage drop along the local circuit, the smart plugs should be installed as close as possible to the PCC. Additionally, smart plugs are designed to monitor only a single phase of the electrical supply.

The use of widely accessible smart plugs for monitoring the distribution grid is primarily constrained by privacy concerns, access to the communication infrastructure, and the accuracy of the measurements, particularly when compared to those obtained from calibrated smart metering systems or power analyzers. Associating the measurement data from individual smart plugs with a specific distribution grid customer could disclose daily routines or installed appliances. Positioning the smart plugs away from the PCC may compromise the validity of the measurements, and the inherent offset of each device must be identified and corrected. Nevertheless, the high frequency of measurements could allow for the mitigation of these limitations, such as through the application of filters. Another practical challenge in field implementation is ensuring access to a reliable Wi-Fi connection for transmitting measurement data. Offering customers the ability to use the smart plug's switching functionality might incentivize them to permit access to their private Wi-Fi networks. From a consumer's standpoint, the configuration process of the device and its functionality would be analogous to that of commercially available smart home devices.

Installing custom firmware on smart plugs for distribution grid monitoring generally voids their warranty. To avoid this, manufacturers would need to offer a software interface

for data collection or enable the smart plugs to connect to custom servers without requiring firmware modifications. Without such support, DSOs would have to flash custom firmware onto the devices before distributing them to customers. The smart plugs evaluated in this study can be remotely flashed without physical access, streamlining scalability. Given the low data transfer requirements, thousands of smart plugs can connect to a single server. If the number of installed devices increases significantly, the workload can be distributed across multiple servers.

With these results, the first research question can be answered: **To what extent can smart plugs be used for distribution grid monitoring?** Similar to related grid monitoring devices, smart plugs can be used for distribution grid monitoring. They have a higher temporal resolution than typical smart meters. However, smart plugs must be installed close to the PCC to provide accurate measurements.

The proposed method for generating pseudo-measurements allows for comprehensive monitoring of the distribution grid with minimal measuring devices. As shown in Table 3.2, simply substituting missing values with the average of measured values yields suboptimal results. By exploiting the homophilic properties of neighboring nodes, known features can be propagated through the graph representing the distribution grid. This approach is computationally efficient, as the features are propagated by multiplying matrices of size  $n \times n$ . In the evaluated test cases, feature propagation outperforms simple averaging of the available measurements, reconstruction using the GINN algorithm, and feature generation using the method proposed in [GFH24b]. The publication by Grafenhorst et al. [GFH24a] outlines the first application of feature propagation in electricity grids. The evaluation demonstrates that this method can reconstruct missing measurements with high accuracy, achieving an MSE of 0.431 p.u. (equivalent to 0.099 V) when 95 % of measurements are unavailable. This surpasses the GINN algorithm by Spinelli et al. [SSU20], which yields an MSE of 0.640 p.u. (or 0.147 V) under similar conditions. This answers the second research question: **How viable is the use of a feature propagation algorithm for pseudo-measurement generation in distribution grids compared to other related methods?**

## 3.6 Conclusion and Outlook

The increasing penetration of distributed renewable generation and new loads, such as electric vehicle charging infrastructure, places unprecedented stress on low-voltage distribution grids. This creates an urgent need for cost-effective, high-resolution monitoring techniques to proactively identify and mitigate grid constraint violations. This chapter evaluates the viability of using ubiquitous, consumer-grade smart plugs as the foundation for such a monitoring system.

First, the measurement accuracy of the hardware is evaluated. By developing a custom, open-source firmware modification, the voltage measurement accuracy of commercially available smart plugs is significantly enhanced, achieving a standard deviation of 0.27 V when benchmarked against a professional-grade power analyzer. The methodology leverages inexpensive smart home hardware that does not require a professional installation and is connected via a WiFi access point to a standard time-series database.

To address the inherent data sparsity of the deployment of such devices, a novel feature propagation algorithm is proposed for generating pseudo-measurements at unmetered nodes. This computationally efficient algorithm, based on the principles of heat diffusion, outperforms both a previously published method [GFH24b] and a state-of-the-art feature imputation algorithm [SSU20] in accuracy while demanding substantially lower computational complexity.

The combination of validated, low-cost hardware and a lightweight algorithm enables real-time grid monitoring. This facilitates the identification of impending congestion and allows for targeted scheduling of DERs to provide ancillary services based on local measurements, enhancing the electricity grid's capability in hosting renewable generation systems and high-power appliances.

Future work will focus on validating the scalability of the feature propagation algorithm on larger, more complex grid topologies with diverse load profiles. Additionally, the propagation of the phase angle  $\theta$  will also be evaluated. Furthermore, the underlying principle of leveraging electricity grid homophily will be explored as a method to accelerate solutions for other critical grid optimization problems, such as the Optimal Power Flow (OPF).



# 4 Optimization of Multi-Energy DERs

---

Distribution grid voltage monitoring is a prerequisite for deploying congestion mitigation strategies. Upcoming congestion scenarios can be detected in advance, and grid-supporting Distributed Energy Resources (DERs) can be scheduled to alleviate such issues. However, determining the optimal sizing and operational schedule of a multi-energy DER is a complex, non-trivial optimization problem. The existence of multiple, competing optimization strategies addressing this problem motivates the third research question of this work: **What are the advantages and disadvantages of optimizing a multi-energy DER using an evolutionary algorithm, an MILP, and an MIQP?** This chapter presents a comparative analysis of these methods to identify the most effective approach for scheduling DERs.

The chapter is based on the publications “Heuristic vs. Analytical Energy Hub Optimization: Design, Implementation, and Trade-Offs” [Gra+24b] and “Grid Aware Portfolio Optimization of a Multi-Energy DER” [GFH24a].

## 4.1 Introduction

Multi-energy DERs can be operated to achieve various objectives. At its core, a multi-objective function is minimized. This function typically encapsulates not only monetary costs but can also incorporate penalties for greenhouse gas emissions or deviations from a target power curve mandated by the DSO. The target power can be based on the real-time grid monitoring and congestion management, such that a DER can support the grid by purposefully adjusting power demand or power injection. Furthermore, a complete analysis of the costs of a DER requires accounting for the capital expenditure during its planning and installation. Therefore, an economically optimal DER operation must also incorporate asset depreciation, which is dependent on the asset size. On the other hand, this relationship allows for the isolated minimization of operating costs, excluding depreciation, by assuming fixed component sizes.

Suitability of an optimization algorithm is assessed against several performance criteria. A primary criterion is computational speed. Load and generation forecasts, as well as grid state estimations, are frequently updated. To continuously recompute the optimal schedule using the most recent data, the algorithm needs to execute at the same cadence as forecast and state-estimation updates. Consequently, minimizing execution time is essential to enabling the timely implementation of the updated schedule.

The algorithm's adaptability is a second critical requirement, given the wide variability of use-cases for multi-energy DERs and diverse component portfolios. Consequently, the framework must facilitate the integration of novel components and the modification of existing ones.

Finally, the quality of the optimization is an essential criterion. An optimal operational schedule for a multi-energy DER consists of operating points for each component across all time steps while pursuing one or more predefined objectives. Objectives include minimizing carbon emissions, enhancing local grid stability, reducing operational expenditures, or a combination of those. The optimization algorithm must identify the optimal solution to the scheduling problem, irrespective of whether the problem consists of a single or multiple objectives. The evaluation of optimization quality can be conducted quantitatively by examining error metrics, such as the MSE.

An optimization algorithm that meets these criteria enables a grid operator to effectively manage the multi-energy DER. The accurate tracking of the target power curve and the timely implementation of the optimized schedule form the foundation for the DSO coordination mechanism proposed in [Ner+25]. The core contribution of this publication is a structured, iterative negotiation between the multi-energy DER optimization algorithm and the DSO. While the interface to the DSO is possible from a modeling standpoint, it needs to be developed as the next practical step. The envisioned exchange begins when the optimization algorithm publishes an initial schedule based solely on the objective of minimizing operational costs of the DER. Subsequently, the DSO can analyze this schedule and request specific changes to mitigate grid congestion. In response, the algorithm re-solves the optimization problem with the new constraints imposed by the DSO. Crucially, the system then computes and publishes the cost difference between the original minimized-cost schedule and the newly requested schedule. This provides the DSO with a transparent, monetized trade-off for its request. Based on this information, the DSO is able to either accept the new schedule and its associated cost—effectively procuring a flexibility service, or request other changes, initiating a new cycle of negotiation. The framework moves beyond simple dispatch commands, creating an economically grounded dialogue for managing grid constraints [Ner+25].

The primary focus of this chapter is a comparative analysis of two optimization algorithms - a heuristic EA and an analytical MILP or MIQP. The algorithms are tasked with the operational scheduling of a multi-energy DER, as this optimization problem is already implemented with an EA. Furthermore, the scheduling algorithm can serve as a foundation for a more comprehensive sizing and planning tool for multi-energy DERs. The framework can be easily advanced to incorporate capital expenditure, such as depreciation. However, to facilitate a direct and equitable comparison between the heuristic and analytical approaches, the scope of this investigation is intentionally confined to a specific multi-energy DER. Consequently, only the operational scheduling aspect is considered.

This chapter continues with a discussion of work related to various techniques to optimize the schedule and design of DERs. A summary of DERs and the Energy Hub can be found in Section 2.1 and is not part of this chapter. The subsequent methods section describes the DER model, which is implemented for an MILP, an MIQP, and an EA. To evaluate these two approaches, the optimization techniques are compared. In the evaluated scenario, the optimization algorithms minimize the difference between a target power curve requested by a DSO and the power curve resulting from scheduling the components of the multi-energy DER. The resulting findings are discussed in the final section of this chapter.

## 4.2 Related Work

Common approaches to scheduling the operation of the components of a multi-energy DER fall into two main categories: data-driven heuristics, such as EAs, and analytical methods, including MILP and MIQP. Indeed, the original Energy Hub by Geidl et al. employs an MILP model [Gei+07]. MILP has since been applied in various contexts, including battery storage scheduling by Chouhan et al. [Cho+16] and by Najafi et al. for a DER model incorporating uncertainty considerations [Naj+22]. Illustrating another analytical technique, Zhong et al. [Zho+22] optimized the schedule for integrated sector-coupling systems using MIQP. The choice of MIQP was driven by their objective function, which, being analogous to an Root Mean Squared Error (RMSE), inherently defines a quadratic optimization problem.

Heuristic optimization techniques, such as EAs [Ela+20] and Particle Swarm Optimization [BAL17], provide an alternative by permitting the direct use of complex, non-linear objective functions and models, thereby obviating the need for linearization [LRM17]. EAs are inspired by the principles of biological evolution. The fundamental process involves iteratively refining a collection of candidate solutions to navigate complex search spaces and identify near-optimal solutions. The core algorithmic loop is defined by a structured sequence as follows.

In the beginning, a population is initialized, representing a collection of candidate solutions. This initial population can be generated randomly to ensure a broad and unbiased exploration of the solution space, or it can be seeded with heuristic-based or educated guesses to potentially accelerate convergence. Each individual solution within this population is formally referred to as a chromosome. A chromosome serves as an encoded representation of a potential solution, and a pre-defined model dictates how this representation can be decoded into a problem-specific solution. Structurally, a chromosome is composed of one or more genes, where each gene encapsulates a set of variables that define a particular aspect or component of the overall solution [BJ09; Pop24].

The iterative cycle of an EA proceeds through the following distinct steps [BJ09; Pop24]:

1. Initial population generation: Initialize the algorithm's population based on heuristics or guesses.
2. Evaluation: Evaluate every individual in the current population by decoding each chromosome and applying a fitness function to quantify its performance.
3. Parent selection: Select a subset of individuals based on their evaluated fitness to act as parents for the subsequent generation.
4. Reproduction: Generate offspring out of the parents through the application of genetic operators, primarily recombination and mutation. Recombination combines the genetic material from two or more parents to create new chromosomes, resulting in

the inheritance of beneficial traits. Mutation introduces small, random perturbations to an offspring’s chromosome, promoting exploration of new regions within the search space.

5. Evaluation of offspring: Evaluate newly generated candidate solutions using the fitness function.
6. Termination Check: Check a set of pre-defined termination criteria, for example, reaching a fixed maximum number of iterations, observing that the population’s fitness has converged (i.e., improvements have stagnated below a certain threshold  $\epsilon$ ), or achieving a pre-determined target value for the goal function.

If the termination criteria are not satisfied, the cycle repeats, starting with step 3. A new set of parents is selected from the current population, and another iteration of reproduction and evaluation is executed. This iterative process continues, progressively guiding the population toward regions of the search space containing higher-quality solutions.

A practical application is the operational optimization of the Energy Hub Gas [Pop+23a], which utilizes the GLEAM evolutionary algorithm [BJ09]. Within this framework, chromosomes produced by the EA are interpreted and translated into Energy Hub schedules. The efficacy of each schedule is then determined through simulations of individual component responses, where operational setpoints are mapped to energy flows (inputs and outputs), and the overall schedule performance is quantified via multiple objective functions.

To decode a chromosome into a schedule for the DER, the individual genes are decoded as described by Poppenborg et al. [Pop+23a] and Jakob et al. [Jak+17]. This gene model describes how the individual decision variables that are encapsulated within a gene are decoded. In this instance, a gene consists of four variables: the component ID, the start time of this setpoint, the duration for which the setpoint is valid, and the setpoint itself, expressed as a fraction of the component’s rated power. Each chromosome consists of 120 to 240 genes, allowing for up to 240 individual setpoints in a one-day schedule of the DER. Five components are scheduled in the Energy Hub Gas [Pop+23a]: a Combined Heat and Power (CHP) plant, a gas storage system, a Proton Exchange Membrane (PEM) electrolyzer, a methanization system, and a Battery Energy Storage System (BESS). A conventional time-discretized formulation would necessitate defining 480 independent setpoints (5 components with 96 time intervals). In contrast, the proposed EA employs a more compact representation. As each gene defines not only a power fraction but also its duration, a single gene can span multiple time intervals. Consequently, a complete 24-hour schedule requires significantly fewer genes to be described, which reduces the necessary computational effort, but possibly also the quality of the solution.

The formulation of analytical optimization problems is commonly done using standalone Algebraic Modeling Languages [Byn+21]. However, modern frameworks such as Pyomo [Byn+21] facilitate this process directly within a high-level programming environment like Python, offering a more integrated workflow [Byn+21]. To facilitate future development of the optimization algorithm, the Pyomo framework is employed for modeling the multi-energy DER.

Although EAs and MILP models are both widely employed for DER scheduling and operational optimization, they possess distinct strengths and weaknesses. The selection of the most appropriate algorithm significantly impacts the quality and feasibility of the resulting DER schedule. Apart from the work "Grid Aware Portfolio Optimization of a Multi-Energy DER" by Grafenhorst et al. [GFH24a], a systematic comparative study of these distinct algorithmic strategies within this context appears to be an unaddressed area in the published research.

## 4.3 Method

In this section, the EA, MILP, and MIQP optimization techniques for a multi-energy DER are outlined. The Energy Hub Gas implementation, as detailed by Poppenborg et al. [Pop+23a], represents the EA optimization.

### 4.3.1 Evolutionary Algorithms

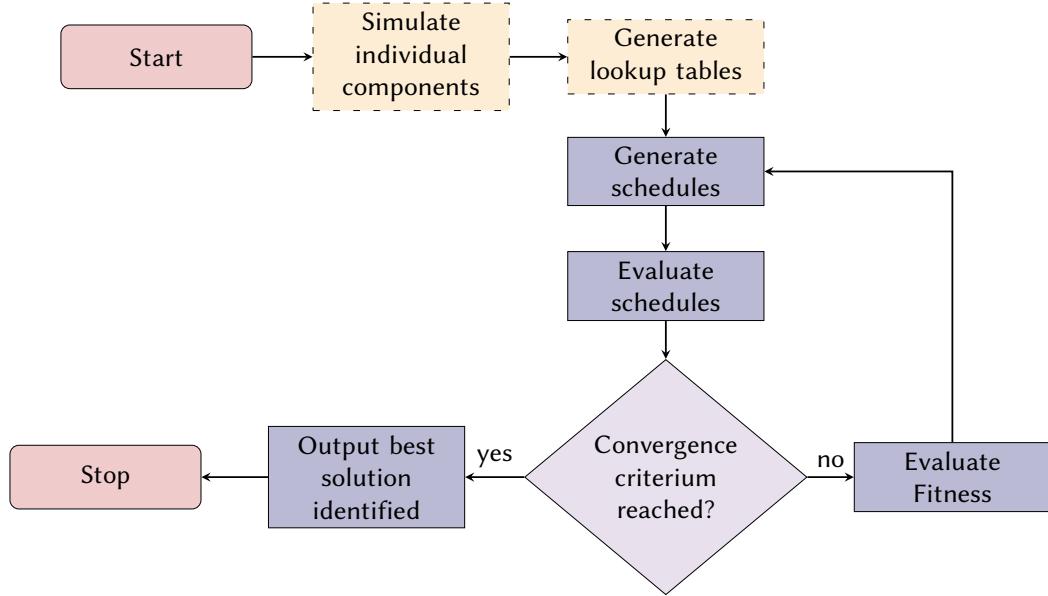
The Energy Hub Gas model [Pop+23a] component representations are based on lookup tables, which delineate the relationship between input and output energy vectors. These tables can be populated either through detailed simulations of individual components or from empirical data obtained from physical components. A key advantage of this data-driven approach is the circumvention of explicit mathematical formulations or model simplifications to characterize these input-output relationships. Consequently, the modeling process becomes largely agnostic to the underlying physical or operational intricacies of the specific components.

The General Learning Evolutionary Algorithm and Method (GLEAM) EA, as utilized in the Energy Hub Gas [Pop+23a], generates candidate solutions. Each solution comprises a series of operational setpoints for all constituent components over a defined time horizon. An Energy Management System (EMS) subsequently evaluates the aggregate energy inputs and outputs of the entire DER by referencing these pre-computed lookup tables for each component. This evaluation enforces operational constraints, including battery State of Charge (SoC) limits, maximum CHP power generation, and maximum PEM electrolyzer load. The process of candidate generation and evaluation is repeated until a predefined convergence criterion is reached. Upon termination, the algorithm outputs the best-performing solution identified during the search. For the comparative analysis herein, this Energy Hub Gas workflow is adopted and visually represented in Figure 4.1.

### 4.3.2 Analytical Model

An MILP optimizes a linear objective function subject to linear constraints. For problems involving quadratic relationships, an MIQP formulation can be suitable. Such quadratic formulations become necessary when the objective function incorporates a quadratic error metric, for instance, the MSE. In these scenarios, an MIQP can yield more accurate solutions.

In the MILP formulation of a multi-energy DER, the components are represented by linearized models, wherein the input-output relationships and inter-component interactions are defined by systems of linear equations. This modeling approach employs a generalized class for all storage systems, irrespective of their technology (e.g., hydrogen or battery storage). Input power is denoted as  $P_{in}$ , and the corresponding energy as  $E_{in} = P_{in} \cdot t$ . Output power and energy are defined analogously. This generalized storage system model



**Figure 4.1:** The EA optimization process used for comparison. The process is based on [Pop+23a]. The single-component simulation and the generation of the lookup tables only need to happen once. The illustration is adapted from [Gra+24b].

is formally presented by

$$soc(t+1) = soc(t) + \frac{E_{in} \cdot \eta_{in} - E_{out} \cdot \eta_{out}}{cap} \quad (4.1a)$$

$$0 \leq P_{in} \leq P_{in,max} \quad (4.1b)$$

$$0 \leq P_{out} \leq P_{out,max} \quad (4.1c)$$

$$0 \leq E_{in} \leq cap - soc(t) \cdot cap \quad (4.1d)$$

$$0 \leq E_{out} \leq soc(t) \cdot cap. \quad (4.1e)$$

To calculate the state of charge  $soc$  of the next time step  $t + 1$ , the  $soc$  of the current time step is modified by adding the energy input  $E_{in}$  multiplied by the charging efficiency  $\eta_{in}$  and subtracting the energy output  $E_{out}$  multiplied by a discharge efficiency coefficient  $\eta_{out}$ .

Constraints (4.1b) and (4.1c) limit the maximum charge and discharge rate of the storage system. Constraint (4.1e) prevents the use of more energy than is stored, and the constraint (4.1d) limits the stored energy to the maximum capacity.

Since the DER is connected to multiple energy grids, the individual connections to the grids are modeled by

$$rev(t) = (E_{out} - E_{in}) \cdot c(t) \quad (4.2a)$$

$$0 \leq P_{in} \leq P_{in,max} \quad (4.2b)$$

$$0 \leq P_{out} \leq P_{out,max}. \quad (4.2c)$$

To calculate the revenue that is generated by selling and importing energy, in equation (4.2a) the difference between the exported energy  $E_{out}$  and imported energy  $E_{in}$  is multiplied by the energy cost  $c(t)$  at every time step. The cost of importing energy is assumed to be equal to the revenue from selling energy. The constraints (4.2a) and (4.2b) limit the maximum power for importing and exporting in a single time step.

$$rev(t) = E_{out} \cdot c(t) \quad (4.3a)$$

$$0 \leq P_{out} \leq P_{out,max} \quad (4.3b)$$

The modeling approach for generation components mirrors that of the grid connection. A simplifying characteristic, however, is their inability to import energy; this unidirectional constraint results in a more concise set of descriptive Equations (4.3).

$$-\sum_{a \in P_{in}} a(t) = \sum_{b \in P_{out}} b(t) \cdot \beta_b \quad (4.4a)$$

$$\forall a \in P_{in} : P_{a,min} \leq a(t) \leq P_{a,max} \quad (4.4b)$$

$$\forall b \in P_{out} : P_{b,min} \leq b(t) \leq P_{b,max} \quad (4.4c)$$

$$\forall b \in P_{out} : \gamma_b(t) = \frac{b(t)}{P_{b,max}} \quad (4.4d)$$

$$\forall b \in P_{out} : \gamma_b(t+1) - \gamma_b(t) < R_{b,up} \quad (4.4e)$$

$$\forall b \in P_{out} : \gamma_b(t) - \gamma_b(t+1) < R_{b,down} \quad (4.4f)$$

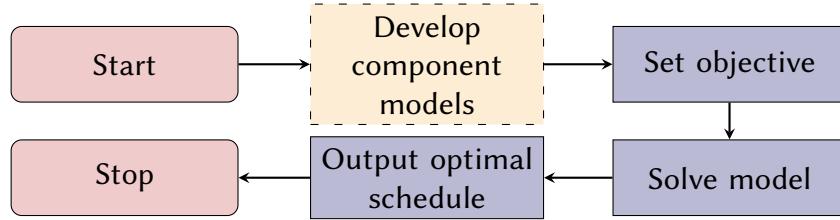
With Equation (4.4a), the relation between the total energy input as a sum of all forms of energy inputs  $a \in P_{in}$  and the total energy output  $b \in P_{out}$  is described. To quantify the losses introduced by the energy conversion processes, the coefficient  $\beta$  is introduced, representing the efficiency of the specific energy output. For each form of energy, constraints (4.4b) and (4.4c) define the minimum and maximum input and output power and equation (4.4d) defines the setpoint  $\gamma$  of each converter component. This setpoint can only be changed by a maximum ramp-up  $R_{b,up}$  or ramp-down  $R_{b,down}$  value as expressed by (4.4e) and (4.4f).

A set of three complementary objective functions is implemented that can be selectively combined. To minimize the deviation from the grid operator's target power profile  $P_t(t)$ , Equation (4.5a) calculates the sum over all time steps of the absolute difference between the target power and the DER power. The economic objective, specified in Equation (4.5b), aims to minimize the operational cost, calculated as the difference between energy input expenditures ( $C_a(t)$ ) and energy sales revenues ( $R_b(t)$ ). A final objective, presented in Equation (4.5c), enables following a target power curve by penalizing deviations according to the RMSE. Consequently, the introduction of this quadratic term reclassifies the optimization problem, converting it from an MILP to an MIQP.

$$\min \frac{1}{T} \sum_{t \in T} \text{abs}(P_{in}(t) - P_{out}(t) - P_t(t)) \quad (4.5a)$$

$$\min \sum_{t \in T} \left( \sum_{a \in P_{in}} a(t) \cdot C_a(t) \right) - \left( \sum_{b \in P_{out}} b(t) \cdot R_b(t) \right) \quad (4.5b)$$

$$\min \sqrt{\frac{1}{T} \sum_{t \in T} (P_{in}(t) - P_{out}(t) - P_t(t))^2} \quad (4.5c)$$



**Figure 4.2:** The MILP optimization process. Development of the component models only has to be done once. The illustration is adapted from [Gra+24b].

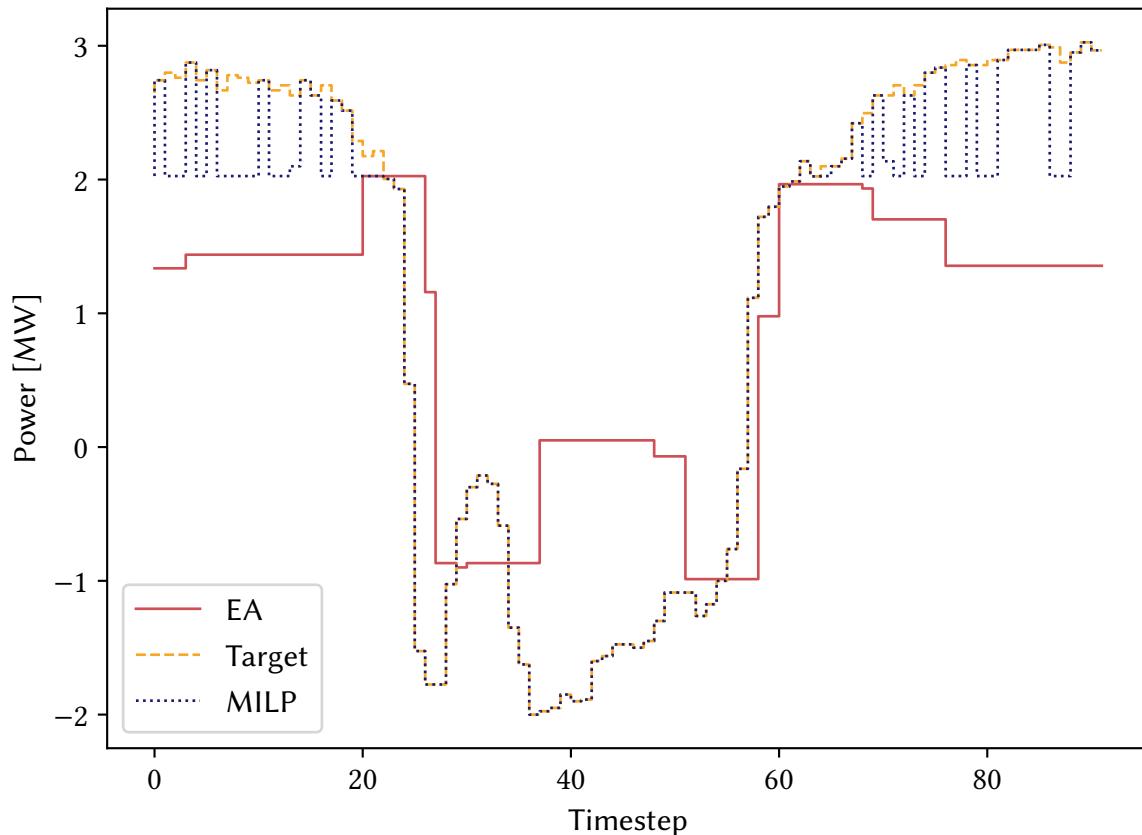
Figure 4.2 outlines the optimization process using the MILP model. First, linear equations representing component behavior are derived from either simulation or empirical evaluation. Next, component-level constraints are formally imposed. The resulting constrained optimization problem is then solved using external computational solvers. The model is implemented in Python environment using the Pyomo library [Byn+21; HWW11].

## 4.4 Evaluation

### 4.4.1 Optimization Results

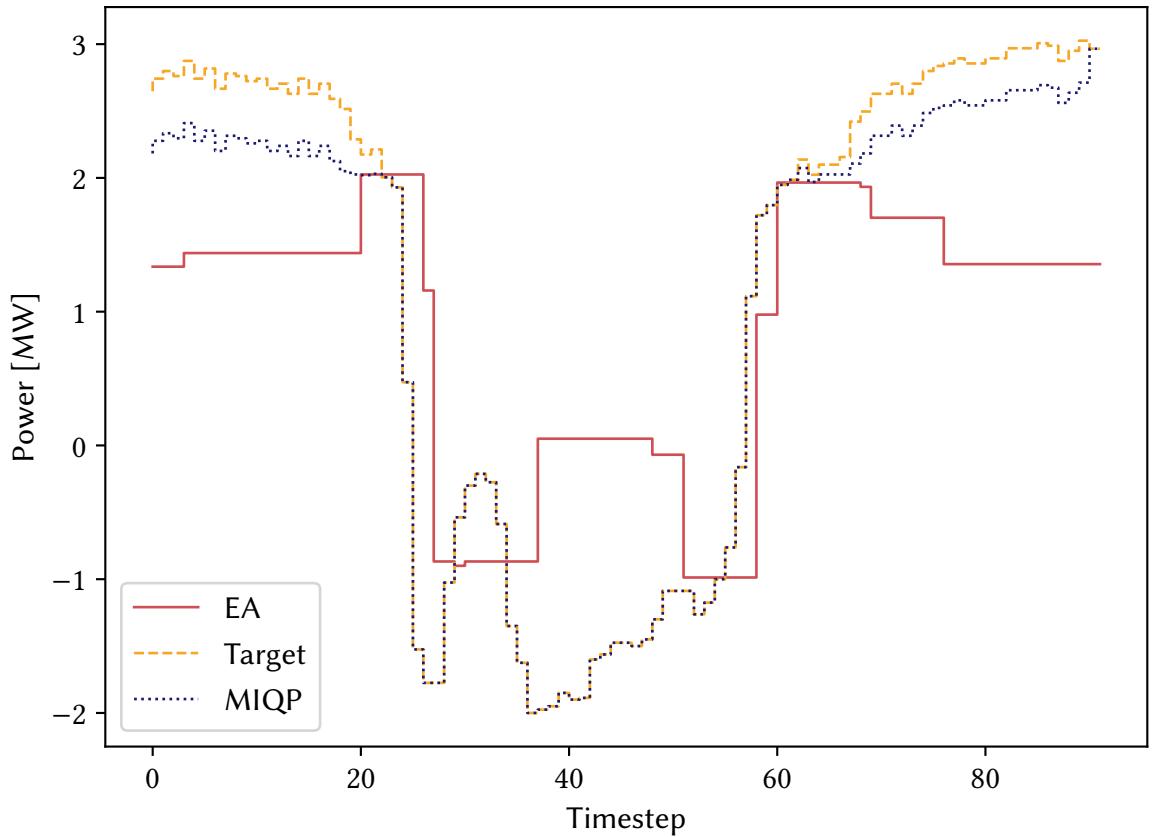
A comparative performance assessment of the MILP and EA optimization methods is performed, wherein both algorithms are tasked with the identical objective of tracking a predefined power target according to the functions in Equations (4.5a) and (4.5c).

The performance evaluation includes a case study for two exemplary days in April 2022, as well as a comprehensive analysis of the entire year 2022. The use of real-world energy prices and solar irradiation data ensures a realistic test environment. As a standard initial condition, all daily simulations start with the battery state of charge at 50 %.



**Figure 4.3:** Comparison of the generated schedules from the MILP and the EA for one exemplary day in April 2022. The Figure is adapted from [Gra+24b].

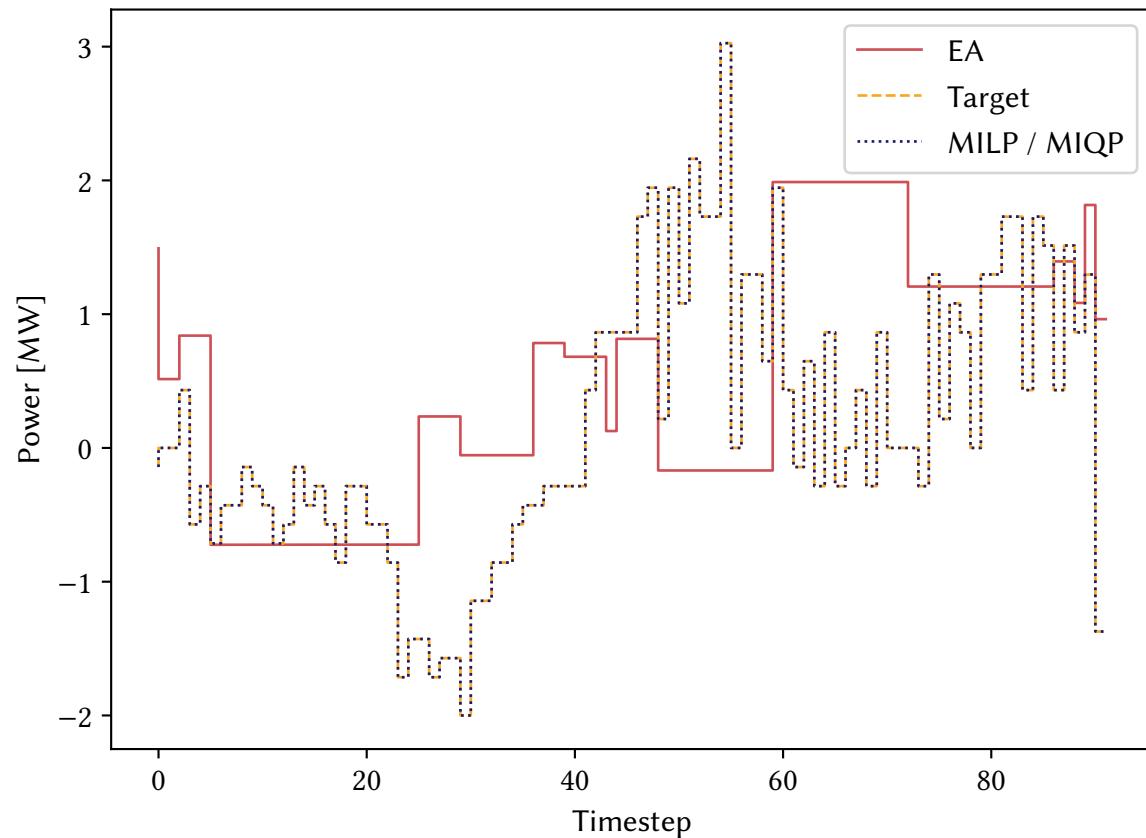
In times without enough generation to match the DSOs target power curve, the schedule optimized with the MILP using the linear error metric (Mean Absolute Error (MAE)) varies between exactly matching the target and differing by a large amount. This is due to the battery storage sometimes supplying enough power to match the target, and at other times not supplying any power at all.



**Figure 4.4:** Comparison of the generated schedules from the MIQP, and the EA for the same exemplary day as is used to generate Figure 4.3. The Figure is adapted from [Gra+24b].

In contrast to the MILP, which minimizes the linear absolute error, the MIQP optimizes the schedule using the RMSE objective function denoted in Equation (4.5c). Consequently, the schedule generated by the MIQP keeps a constant deviation from the target curve in the times with insufficient generation capacity at the start and the end of the day. During these times, the BESS is being discharged at a constant rate. Doing so minimizes the sum of the quadratic error, while the sum of the absolute error as calculated with Equation (4.5a) is

equivalent to the solution generated using the MILP.



**Figure 4.5:** Comparison of the generated schedules from the MILP, the MIQP, and the EA for a separate exemplary day. In this case, the results of the MILP and the MIQP are identical and match the target exactly.

Figure 4.5 illustrates the optimized DER schedule for another day. In this specific scenario, both the MILP and MIQP algorithms result in an identical component schedule. This result is attributable to the sufficient capacity of the DER components, which allows the target power curve to be tracked perfectly. Notably, the EA maintains its scheduling performance observed in Figure 4.3, continuing to employ a limited number of setpoint adjustments rather than adjusting the DER's power at every time step.

While the MILP formulation penalizes deviations linearly using the absolute difference (Equation (4.5a)), the EA and MIQP employ the RMSE (Equation (4.5c)). This quadratic objective function disproportionately penalizes larger deviations. This fundamental difference in the objective function explains why the MIQP consistently yields schedules with a lower or the same RMSE compared to the MILP optimization, a conclusion substantiated by the comprehensive statistical analysis for an entire year, as documented in Table 4.1.

RMSE (MWh)	EA	MILP	MIQP
Average	0.97	0.21	0.14
95th percentile	1.18	0.46	0.34
Standard Deviation	0.14	0.16	0.12

**Table 4.1:** Results of optimizations for every day of the year 2022. The RMSE is calculated for every day in one year. The data was first published in [Gra+24b].

The EA is a heuristic algorithm that terminates upon meeting a convergence criterion (Figure 4.1). Consequently, it provides an approximation of the optimal solution rather than the exact optimum. Furthermore, the optimization's precision is constrained by the use of discretized lookup tables for component modeling, which, due to their finite number of entries, do not offer an exact representation of component behavior.

An observational difference is that the EA produces schedules with fewer setpoint changes compared to the MILP and the MIQP. This behavior is a direct consequence of the solution encoding scheme. As described by Poppenborg et al. [Pop+23a] and Jakob et al. [Jak+17], each chromosome contains decision variables encapsulated in genes that represent the DER's component, start time, duration, and power fraction. However, this chromosome structure is limited by the number of genes and cannot account for individual power settings for each component in each time interval.

Figure 4.5 compares the 96-timestep schedules optimized by the three algorithms for the DER. Among the methods, the schedules generated by the MILP and MIQP demonstrate the highest degree of accuracy in following the specified target power profile.

#### 4.4.2 Runtime

The runtime comparison focuses solely on the optimization task, as the one-time model generation cost is excluded from the analysis of the daily DER scheduling. For a 96-timestep daily schedule, the MILP model requires 1.6 seconds to solve with GUROBI [Gur23] on a modern office computer, while the MIQP requires 2.0 seconds. In contrast, the employed EA, which must evaluate many potential schedules, has a significantly higher runtime of approximately 20 minutes.

#### 4.4.3 Configurability

The EA utilizes lookup tables populated by Functional Mock-up Unit (FMU) simulations. This methodology enables straightforward component integration, requiring only the simulation of a new component's FMU.

The MILP model for DER optimization is manually formulated by expressing all components and their interactions as linear equations. While this initial linearization is time-consuming, the model architecture supports future expansion. New components can be added by either creating new linear representations or by combining the models of existing components.

The MILP and MIQP models are open-source. This allows other researchers to use them as a foundation, avoiding the time-intensive process of creating similar models from scratch. Furthermore, its architecture supports the integration of other open-access models, such as those in [Sas+20].

## 4.5 Discussion

The results of this study demonstrate the practical utility of the MILP modeling approach for the operational scheduling of DERs and their coordination with a DSO. The MILP algorithm computes the exact, globally optimal solution for its underlying linear model. However, because the model itself is an approximation, this optimality does not extend to the real multi-energy DER. Inaccuracies are introduced during the required linearization of system components (Figure 4.2). Therefore, while the algorithm provides a calculated, optimal result, this result is subject to the modeling errors introduced by the linearization of the components and objective functions.

As outlined in the previous section, the EA produces schedules with a low temporal resolution. This phenomenon is a direct consequence of the solution encoding, wherein a single gene represents a constant power output over multiple time intervals. In contrast, the target power curve fluctuates, hindering precise tracking of the target power curve. To address this limitation, Poppenborg et al. develop a more sophisticated two-stage dynamic optimization algorithm [Pop+23b], based on the original Energy Hub Gas EA. Their approach performs an initial optimization run, followed by a second, targeted run that focuses on refining the schedule during periods of high tracking error. While this method yields improved accuracy, it does so at the cost of a substantial increase in computational effort. With the applied enhancements, the EA-based method exhibits improved accuracy but still falls short of the tracking performance of MIQP. This implies that alternative, untested EA variants could potentially outperform MIQP in accuracy. Nevertheless, the findings of this chapter indicate that EAs are unlikely to exceed MIQP in computational efficiency, at least for multi-energy DER optimization with a similar number of linear equations.

All evaluated algorithms can be extended in order to support new components. For the EA, a simulation or real-world evaluation of the component is necessary. To generate the required look-up table, it is necessary to test all conceivable power set-points, along with their respective input and output powers. However, the configuration of the EA is a complex challenge that necessitates either a dedicated hyperparameter search or the adoption of parameters from existing literature [BJ09]. These hyperparameters include the chromosome length, initialization of the genes, offspring selection method, crossover and mutation functions, and evaluation function. The integration of novel components may lead to new local and global minima. Consequently, the hyperparameters must be adjusted to ensure that the EA finds the global minimum. Since Poppenborg et al. extensively optimized hyperparameters for the multi-energy DER use case [Pop+23a], their approach is taken as representative of the broader family of EAs. Adapting another EA to this use case is beyond the scope of this work, as the required hyperparameter tuning would be too resource-intensive.

For the MILP and MIQP, new components can be added by linearizing the components' behavior. The effort required for this process is contingent upon the specific component in question. However, the developed model's modular and object-oriented design enables novel components to inherit the generic functionality of an energy storage system, an energy conversion system, an energy generation system, or a generic load.

Given that the DER scheduling problem is well-defined and the constituent component models can be linearized with negligible loss of fidelity, an MILP model provides a natural and appropriate optimization approach. For the test case, the MILP solver outperformed an EA, yielding superior results in significantly less time. The model can also be extended to an MIQP formulation to directly optimize quadratic objectives like the RMSE. These findings answer the third research question: **What are the advantages and disadvantages of optimizing a multi-energy DER using an evolutionary algorithm, an MILP, and an MIQP?**

## 4.6 Conclusion and Outlook

This chapter investigates two distinct approaches to scheduling a multi-energy DER: a heuristic method, represented by an EA, and an analytical method, represented by MILP and MIQP models. The EA's models are built using a co-simulation of a DER. Consequently, the optimization is restricted to operational scheduling, and the MILP cost function likewise omits depreciation. The behavior of the co-simulation components is individually approximated through linear functions, as described in Equations (4.1), (4.2), (4.3), (4.4), and (4.5).

While the presented evaluation focuses on two exemplary days, the results are generalizable. Both algorithms consistently exhibit comparable performance profiles and distinguishing characteristics across all tested days. This is exemplified in Figure 4.5, which illustrates a day where the MILP and MIQP formulations result in an identical, optimal schedule. In this scenario, the EA again displays its characteristic behavior of infrequent setpoint adjustments. This results in lower performance compared to the MILP and MIQP optimizations, because the fluctuating target power can not be met accurately.

Two primary research directions are identified to advance the optimization of DER scheduling. First, the process of manually formulating mathematical models presents a bottleneck. Future work should focus on the development of automated modeling tools that can generate the MILP or MIQP representations from a high-level component description or simulated components. This would not only substantially reduce development effort but also enhance the scalability of the optimization approach. Second, extended benchmarks are required to evaluate the performance of the optimization algorithms across varying problem sizes. It is hypothesized that for large-scale problem instances, the computational efficiency of an EA may surpass that of a deterministic MILP solver. However, the problem size, structure, and parameterization at which an EA becomes more effective than an MILP has not yet been investigated and is subject to future research.

Planning and sizing of a DER must take into account the physical limitations of its hosting energy grids. This is particularly important for constrained energy systems that are operating close to or at their limits. Therefore, this work develops an integrated approach, explicitly modeling the constraints imposed by the surrounding electricity grid within the optimization framework for the design of the DER. This chapter addresses configuring and sizing a DER, extending the operational scheduling optimization from the previous chapter. The primary objective remains economic optimization, aiming to minimize the total operating cost, including asset depreciation. In alignment with the findings of the preceding analysis, an MILP/MIQP model is utilized. The key distinction of this advanced model is the inclusion of capital expenditure. This is achieved by augmenting the linear model to include component depreciation, thereby creating a model that balances initial investment against long-term operational savings. By implementing and evaluating this model, the fourth research question is answered: **How does the annual cost of a DER designed to provide grid supporting services compare to the annual cost of the most economical design?**

This chapter is based on the publication “Grid Aware Portfolio Optimization of a Multi-Energy DER” [Gra+24a].

## 5.1 Introduction

As detailed in the previous chapter, the schedule of a DER can follow a target power curve. By demanding or injecting power according to the target power curve, the grid is supported. To avoid excessive increases in the cost of operation of the DER, the deviation from the most economical power curve should be minimized. To achieve this, the target power curve should only constrain the DERs power at times when it is necessary. Respecting such constraints is also necessary during the planning and sizing of the individual components of a DER to achieve the cost-optimal configuration for a multi-energy DER that can fulfill the grid’s needs. Calculating these constraints requires an analysis of the grid area surrounding the DER in conjunction with optimizing the composition of a DER.

This work proposes an algorithm to determine the optimal economic sizing of DER components while ensuring the sufficient provision of flexibility to mitigate grid congestion. The methodology integrates the DER optimization presented in the previous chapter with power system analysis, necessitating complete grid monitoring. As the optimal placement of DERs constitutes a high-dimensional, non-convex problem [GGK18], it is deliberately excluded from the scope of the present analysis. Consequently, the multi-energy DER is

assumed to be situated at a fixed location. In practice, the locations that are feasible for the installation of a DER are often governed by non-electrical factors, such as the availability of land and points of grid interconnection. This lends credence to a fixed-location analysis. The process unfolds as follows: first, power flow calculations identify grid congestion events and determine the requisite active power adjustment to restore normal operation. Second, these requirements constrain an optimization problem that co-optimizes the sizing of the DER components and their operational schedule. An iterative procedure is essential due to the inter-temporal dependencies; shifting a load from one time step to another necessitates a re-evaluation of the grid state and the optimal schedule across the entire horizon. A central contribution of this work is the model's capacity to directly incorporate the operational needs of the surrounding electrical grid, enabling the systematic design of cost-effective, grid-enhancing energy systems. The optimization framework utilizes realistic input data, such as historical German electricity prices from 2019 and corresponding PV generation profiles, to inform the analysis. Specifically, the model can be constrained by time-series data defining the required minimum and maximum power generation or load at each time step. This ensures that the final optimized DER design is explicitly provisioned to fulfill these predefined grid support functions. The total costs calculated within the presented study represent the asset depreciation and operating costs for a whole year.

The next section reviews work related to grid-dependent optimization, including various approaches that integrate a secondary objective. Afterwards, a method is developed to optimize the composition of a DER for a specific grid area. Based on the findings presented in the previous Chapter 4, an MILP is chosen as the optimization algorithm. In the following case studies, this method is applied to two distinct scenarios. The chapter concludes with a discussion of the results, answering the fourth research question: **What is the most economical design of a DER for a specific grid area?**

## 5.2 Related Work

When planning and sizing a DER, constraints stemming from the surrounding grid need to be modeled. An illustrative example of a decoupled approach is the work of Ghasemi et al., who employ a bi-level decomposition strategy to separate the DER's design optimization from grid expansion planning [Gha+22]. This technique successfully reduces computational complexity at the cost of sacrificing guaranteed global optimality. More fundamentally, by treating the electrical grid as an external element rather than an integral part of the primary optimization, their planning methodology does not account for the operational constraints the grid imposes. Consequently, the resulting DER design and dispatch schedules may not be viable when deployed within capacity-constrained or weak grids [Gha+22].

Existing MILP-based approaches for grid-aware planning often face scalability and accuracy limitations. Mashayekh et al. consider the surrounding grid but limit their scope to small microgrids with less than 10 nodes, employing power flow linearization [Mas+17]. In the domain of gas grids, Shao et al. use piecewise linearization to formulate a single, comprehensive MILP for system planning [Sha+17]. The primary drawback of such monolithic approaches is that for large grids, they become either computationally expensive or necessitate simplifications that degrade solution accuracy. A further constraint is the inflexible requirement to model the entire grid within the MILP framework.

In the domain of grid-integrated energy storage, Grover-Silva et al. present an approach based on an optimal power flow algorithm for the simultaneous optimization of battery placement and sizing [GGK18]. They demonstrate the method's applicability to large distribution networks (137 nodes) and over extended time periods (130 time steps). The central contribution, as declared by the authors, is the unified consideration of both the locational and capacity-planning aspects of battery deployment. While significant, the scope of their optimization is intentionally narrow, being exclusively confined to battery storage systems. Consequently, it does not address the more complex and multifaceted problem of optimizing the composition and sizing of a multi-energy DER.

La Scala et al. formulate a nonlinear, multiobjective optimization problem designed to simultaneously minimize energy costs for meeting demand while ensuring robust and reliable grid operation [LVZ14]. Their framework models a multi-carrier energy grid wherein multiple, distinct components are interconnected. The optimization focuses on controlling these interconnections to manage the energy flows between the systems optimally. Notably, the scope of their approach is confined to the interaction between these systems; it does not extend to the optimization of the internal component composition of the individual energy systems themselves.

Poppenborg et al. concentrate on the detailed modeling of energy hub components, particularly for renewable gas systems, with an EMS determining the optimal schedule [Pop+21]. Key limitations of their approach include a fixed energy hub composition, as the component mix is not subject to optimization. Additionally, a related study assumes the hub is controlled via an external signal from the grid operator, rather than through a real analysis of the surrounding energy grids [Pop+22].

### 5.3 Method

In this context, grid congestion is analogous to a line's utilization exceeding a limit during specific time steps. The line utilization in Pandapower [Thu+18] is defined as  $loading\_percent = \frac{I}{I_{max}} \cdot 100$ . To reduce the line utilization, the DER must possess the necessary load or generation power. This presents a dual optimization challenge: the component composition of the DER must be strategically designed to provide the required power, while its operational schedule must be concurrently optimized to deploy these assets in the most economically efficient manner.

The cost of an asset is assumed to scale linearly with its size, resulting in a constant cost per unit of power for each component within the multi-energy DER. These costs are divided into Capital Expenditures (CAPEXs), covering procurement, and Operational Expenditures (OPEXs), for asset operation. The CAPEX is amortized over the lifetime of the respective asset. To ensure realistic modeling, hourly European electricity price data is utilized<sup>2</sup>. A summary of all cost assumptions is provided in Table 5.1.

Component	$C_{capex}$	$C_{opex}$	Lifetime (years)
PV system	1300 €/kW	-	20
CHP plant	550 €/kW	12 €/kW · a	15
PEM	520 €/kW	12 €/kW · a	20
BESS	350 €/kWh	5 €/kWh · a	10

**Table 5.1:** Cost assumptions for the individual components of the multi-energy DER.

In optimizing the DER's composition for a specific location, a primary objective is to impose minimal a priori restrictions on the solution space while simultaneously ensuring computational efficiency. Two naive approaches to incorporating grid constraints present significant drawbacks. A static, brute-force application of a single, worst-case constraint across all time steps results in an overly conservative and oversized DER design, as it fails to account for temporal variations in grid conditions. A multi-energy DER dimensioned for the worst-case constraint would have to be able to supply the required load or generation for the worst-case scenario at every time step. Conversely, an exhaustive, a priori calculation of precise constraints for every individual time step imposes unnecessary computational expense, rendering the problem intractable for realistic time horizons.

Therefore, an adaptive constraint calculation strategy is developed. Instead of static and exhaustive constraints, the presented approach identifies and incorporates constraints dynamically, only for the specific time steps where operational limits are violated. The iterative procedure that implements this constraint calculation is presented in Algorithm 2.

<sup>2</sup> <https://github.com/EnergieID/entsoe-py>

---

**Algorithm 2** Optimization of the composition of the DER

---

```

1: procedure OPTIMIZE(Grid  $g$ , Bus  $b$ , Line  $l$ , UtilLim  $lim$ )
2:    $T_{crit} = []$  ▷ Critical time steps
3:   repeat
4:     Power[]  $P_{MILP} = \text{SolveMILP}(\text{constraints})$ 
5:      $g(b).\text{setPower}(P_{MILP})$ 
6:      $T_{crit} = []$ 
7:      $l_{util}^T = \text{calcPowerFlow}(g^T)$  ▷  $\forall t \in T$ 
8:     for all  $t \in T$  do
9:       if  $l_{util}^t > lim$  then
10:         $T_{crit}.\text{add}(t)$ 
11:       end if
12:     end for
13:     for all  $t \in T_{crit}$  do
14:        $P_T^t = \text{reduceUtil}(g^t, b, l, lim, t)$ 
15:        $\text{setConstraints}(P_T^t, P_{MILP}^t, t)$ 
16:     end for
17:   until  $T_{crit}.\text{empty}()$ 
18: end procedure

```

---

The iterative process is initialized by computing an optimal configuration and schedule of the multi-energy DER using the MILP solver, initially without considering grid constraints (line 4). This resulting power profile is then applied to the bus  $b$  in the grid (line 5), and a power flow analysis is executed for all time steps (line 7). The subsequent **for all** loop (line 8) iterates the results to identify the set of critical time steps,  $T_{crit}$ , where line utilization limits are violated. For each identified violation, a second loop calculates a new target power  $P_T^t$ , restricting the DER's generation or load. This power target is then incorporated into the MILP as a new constraint (line 15) via the procedure in Algorithm 4. In case the DER's injected power needs to be limited or the load of the DER needs to be increased ( $P_T^t > P_{MILP}^t$ ), a lower power limit is set. Otherwise, when the injected power needs to be increased or the load needs to be lowered, an upper power limit is set. This entire loop is repeated, with the MILP being re-solved in each iteration, until no grid violations are detected, at which point the *criticalTimesteps* list is empty and the algorithm terminates. To calculate the target power ( $P_T^t$ ) needed to reduce the line's utilization at timestep  $t$  below the predefined limit  $lim$ , the following Algorithm 3 is called (line 14 in Algorithm 2).

---

**Algorithm 3** Reduction of the utilization of line  $l$ 

---

```

1: procedure REDUCEUTIL(Grid  $g^t$ , Bus  $b$ , Line  $l$ , UtilLim  $lim$ , TimeStep  $t$ )
2:    $l_{util}^t = \text{calcPowerFlow}(g^t)$                                  $\triangleright$  line utilization
3:    $m = 1$                                                         $\triangleright$  modifier
4:    $\epsilon = 0.1$ 
5:   while  $m > \epsilon$  do
6:      $g^t(b).add(m)$                                                $\triangleright$  modify DER power
7:      $l_{util}^{prev} = l_{util}^t$ 
8:      $l_{util}^t = \text{calcPowerFlow}(g^t)$ 
9:     if  $\text{abs}(l_{util}^t - lim) > \text{abs}(l_{util}^{prev} - lim)$  then
10:       $m = m \cdot -0.5$ 
11:    end if
12:   end while
13:   return  $g^t(b)$ 
14: end procedure

```

---

The Algorithm 4 functions as a simplified optimal power flow where the only variable is the active power at the DER's bus,  $g^t(b)$ . The algorithm iteratively modifies the power value (line 6), calculates the power flow, and checks the new line utilization  $l_{util}^t$ . If the deviation from the line utilization limit increases, the step-size modifier  $m$  is multiplied by  $-0.5$  to reduce the step size and reverse the search direction. This process, which is closely related to a binary search, terminates when the modifier  $m$  is less than a tolerance  $\epsilon$  and the utilization converges to the limit. The resulting power at bus  $b$  becomes the DER's target power  $P_T^t$  for time step  $t$ . The algorithm can be parallelized by creating individual processes for each time step, yielding a linear speedup.

---

**Algorithm 4** Generation of MILP constraints

---

```

1: procedure SETCONSTRAINTS(Power  $P_T^t$ , Power  $P_{start}^t$ , TimeStep  $t$ )
2:   if  $P_T^t > 0$  then
3:     if  $P_{start}^t > P_T^t$  then
4:       setMaxLoad( $t, P_T^t$ )
5:     else
6:       setMaxGen( $t, -1 \cdot P_T^t$ )
7:     end if
8:   else
9:     if  $P_{start}^t > P_T^t$  then
10:      setMaxGen( $t, P_T^t$ )
11:    else
12:      setMaxLoad( $t, -1 \cdot P_T^t$ )
13:    end if
14:  end if
15: end procedure

```

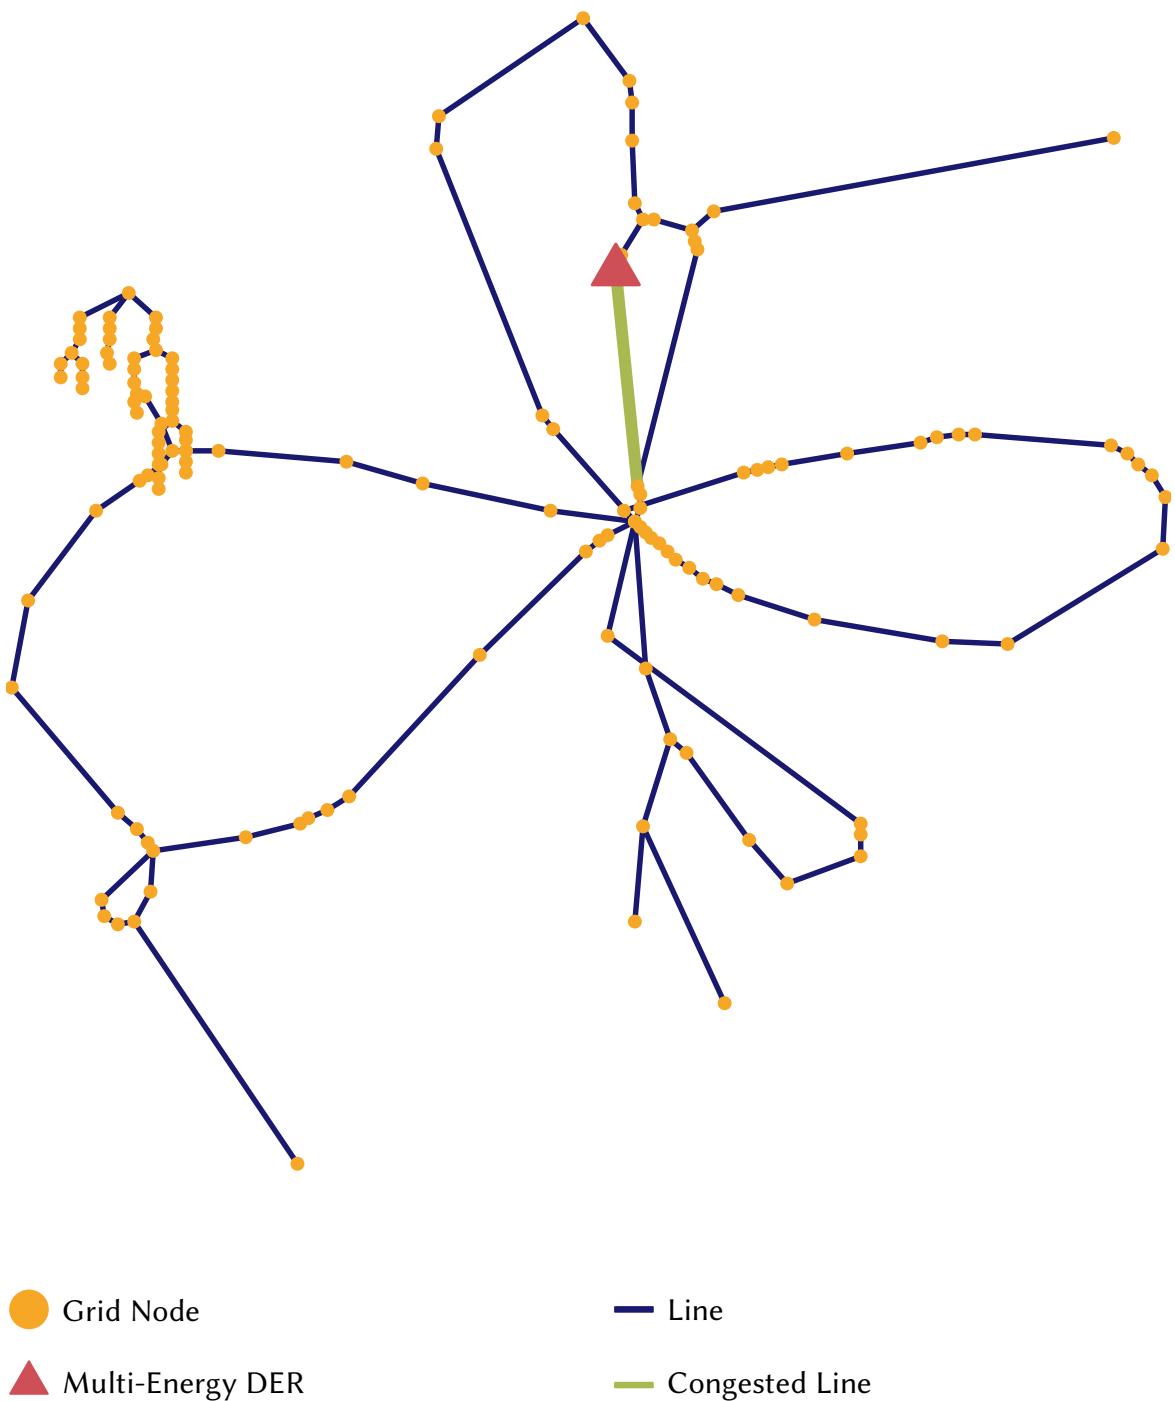
---

The power value calculated by Algorithm 3 corresponds to the approximate operational point at which the line's utilization reaches its maximum permissible value. Consequently, this power value defines the boundary for the DER's power (i.e., its maximum power injection or demand) without violating grid constraints. This operational boundary is then used to formulate a new constraint for the MILP model. The nature of this boundary constraint is determined by comparing it to the power level in the original, unconstrained schedule at the same time step (line 3). This either limits the load or generation and ensures that the new schedule remains within the feasible region.

The algorithm's constraint generation mechanism is best illustrated by an example. Assuming the DER is asked to provide a load and  $P_T^t$  is positive, the unconstrained DER may request more power than is required for reducing the line's utilization. Therefore, the line's utilization is below the maximum permissible value while the DER operates at the most economical setpoint. In this case, the maximum generation is set to the calculated maximum load value (line 4), but the new constraint does not influence the DERs schedule. Conversely, a different scenario arises when the unconstrained DER demand is below the calculated power limit that is needed to reach the line's maximum permissible utilization. In this case, grid conditions necessitate increased consumption, for instance, to absorb surplus local generation and prevent power transmission across highly utilized lines. The DER therefore needs to increase its load. In this case, a lower bound on consumption is enforced, equivalent to limiting negative generation (line 6). This principle of constraint application is mirrored for cases where the DER must inject power and  $P_T^t$  is negative, enabling the algorithm to enforce either an upper or a lower bound on its generation as required by the grid.

## 5.4 Case Studies

The proposed algorithms facilitate the optimization of a DER design that is tailored to the requirements of a local grid area. In this study, line utilization is employed as the primary metric to demonstrate how power flow solutions can be integrated into the MILP model to generate grid-aware constraints. The algorithm is, however, extensible. Other metrics, for example nodal voltage magnitudes, can be readily incorporated with minor modifications to Algorithm 3. A necessary consequence of incorporating these grid constraints is that the total cost of the DER exceeds that of an unconstrained, economically optimal solution.



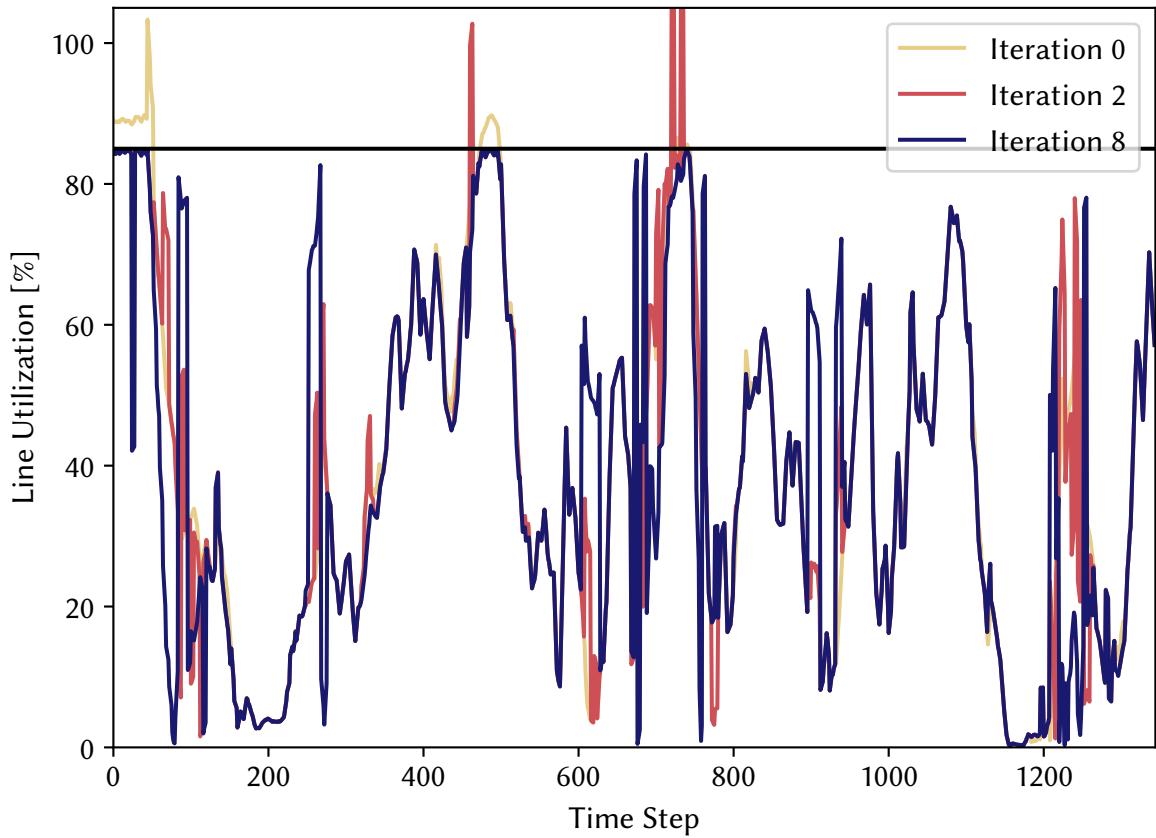
**Figure 5.1:** Illustration of the SimBench grid used in the case studies. The red triangle depicts the location of the multi-energy DER, and the green line depicts the congested line. Orange dots illustrate nodes.

### 5.4.1 Case Study 1: Proof of Concept

The developed method is demonstrated on the 140-node SimBench grid "1-MVLV-rural-4.101". The objective is to reduce high line utilization caused by a wind turbine to below a limit of 85 %. A controllable multi-energy DER shall be located at the same node (Figure 5.1). The initial case study illustrates the iterative optimization process for a DER comprising only a PV system and a battery storage system, formulated as an MILP model. The analysis covers 1344 quarter-hour time steps (14 days). The time frame may be adjusted to accommodate either a full year of operation or a specified number of days. The duration of this time frame is directly proportional to the number of time steps that exceed the line utilization limit. Consequently, the algorithm requires a greater amount of time to compute. The components of the multi-energy DER are dimensioned to meet the calculated requirements of the selected days; thus, choosing representative days across the year results in a multi-energy DER capable of satisfying requirements throughout the entire year.

Upon integrating the cost-optimal multi-energy DER into the grid, the line utilization limit of 85 % is violated at multiple time steps (Figure 5.2). In response, the algorithm generates new constraints specifically for these critical time steps. This intervention successfully reduces the line utilization to approximately 85 % during these critical time steps. Minor deviations from the exact limit are attributable to the numerical approximation inherent in the method. Following the imposition of these constraints, the MILP solver re-optimizes the schedule, reallocating energy to charge the battery in the most economically efficient manner possible while adhering to the newly calculated constraints. A notable outcome of this re-optimization is the shifting of load, which can increase line utilization at previously unconstrained time steps. In response to line utilization limit violations, new constraints are calculated using Algorithm 3 and set via Algorithm 4. However, intermittent violations of the 85 % limit persist, indicating that the algorithm has not yet converged and requires further iterations. After eight steps, the limit of 85 % utilization is no longer exceeded. The initial line utilization, the line utilization after the second, and after the eighth iteration are depicted in Figure 5.2.

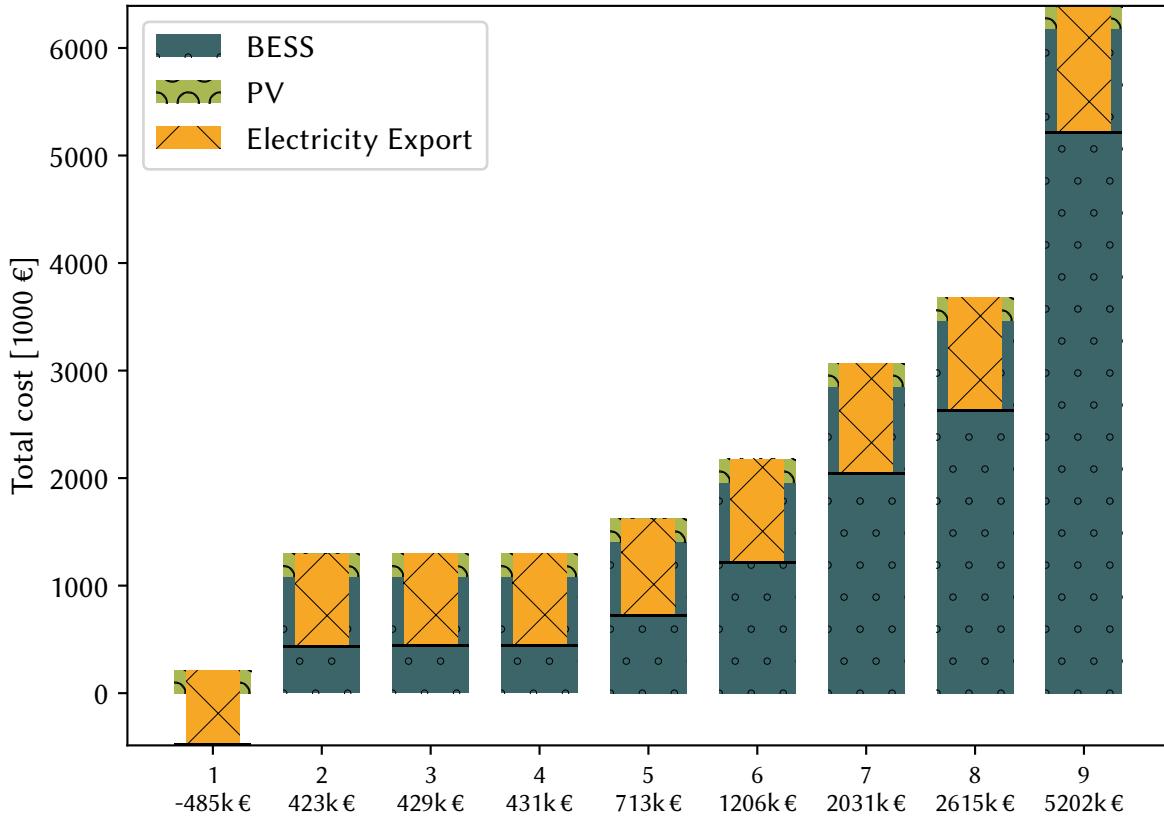
This case study illustrates the iterative optimization process and underscores the necessity of planning and sizing the DER for the specific grid location. The initial, unconstrained optimization yields a multi-energy DER composition that excludes a BESS. However, the target power curve  $P_T^t$  mandated by the DSO requires the DER to function not only as a generator but also as a load at specific times. In the scope of this case study, the only component capable of fulfilling this energy consumption requirement is the BESS. Consequently, as the iterative algorithm incorporates these grid-support constraints, the inclusion and sizing of the BESS become imperative, leading to its progressive increase in capacity in subsequent iterations.



**Figure 5.2:** Line load at the start, after the second, and after the eighth iteration of Algorithm 2. “Each time step represents 15 minutes, with the 1344 time steps shown in the figure corresponding to 14 days.” [Gra+24a]

The optimization algorithm initially provisions the smallest possible BESS capacity necessary to meet operational requirements. This minimal sizing is a direct economic consequence of the energy price structure, where the purchase price of electricity exceeds the selling price, rendering energy arbitrage unprofitable. Nevertheless, even this minimal BESS capacity induces temporal load shifting in the dispatch schedule, as illustrated in Figure 5.2. As the iterative process continues, additional constraints mandate greater operational flexibility, which in turn necessitates a progressive increase in the BESS capacity. This enlargement of the BESS directly corresponds to an increase in the total system cost.

Figure 5.3 depicts the component configuration and total cost for each stage of the optimization. In the initial, unconstrained step, the optimal solution does not include a BESS. The primary economic components are the PV system, which incurs a depreciation cost, and the revenue generated from selling its electrical output. Within the diagram, the cost associated with the PV system is represented by the yellow bar, while the profit from electricity sales is indicated by the green bar. The net result for this initial configuration is



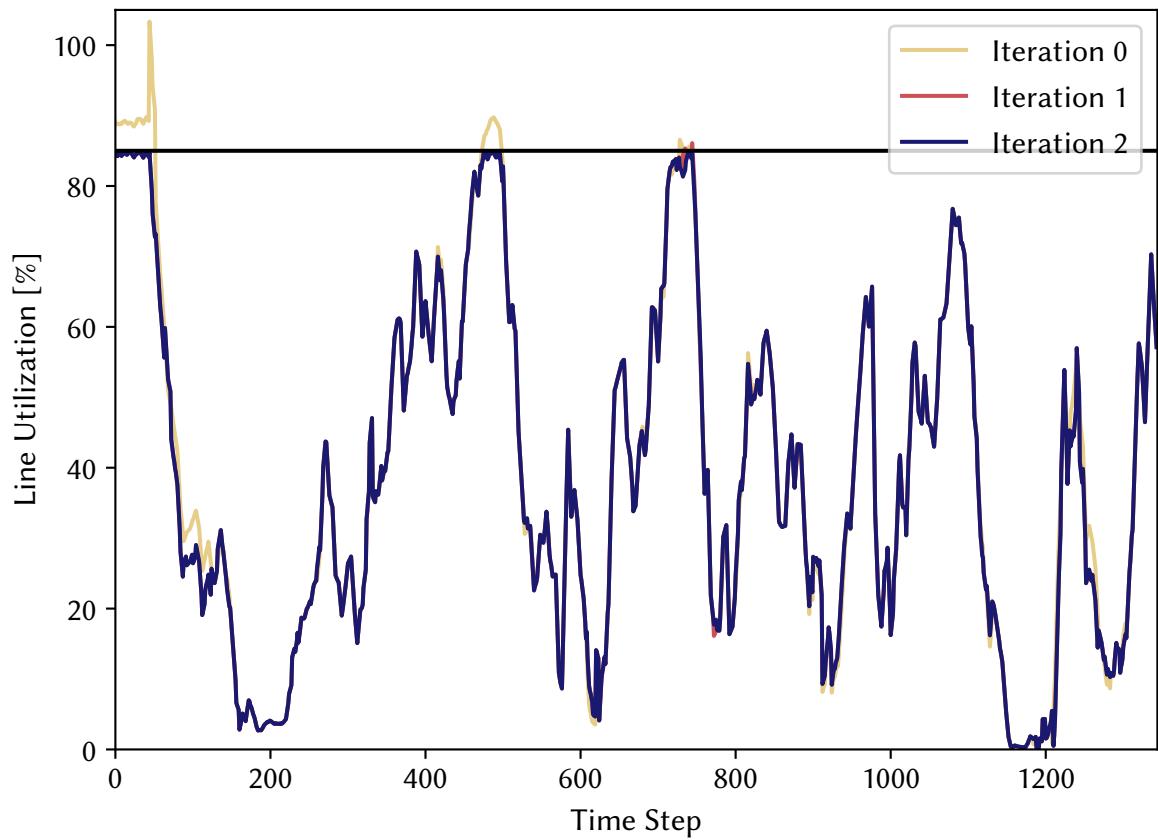
**Figure 5.3:** Illustration of the evolution of the DER composition and its associated costs for a whole year in each iteration of the optimization process. A clear trend emerges: as new grid-aware constraints are introduced in each iteration, the total system cost exhibits a monotonic increase. In the chart, total expenditures are represented by the thicker, upward-extending bars, while revenues from energy sales are depicted as thinner, downward-extending bars.

a total profit (negative cost) of 458k € for an exemplary year, denoted by the black line at the bottom of the green bar. In the subsequent iterations, the total system cost increases in direct proportion to the required BESS capacity. This increasing cost of the BESS is visually represented by the red bar in Figure 5.3. Although the DER continues to generate revenue from electricity sales, this income is insufficient to offset the capital cost of the BESS, resulting in a positive total cost from the first constrained iteration (step 1) onward. The optimization converges to a final configuration with a total cost of 5202k € for an exemplary year. This value represents the minimum economic investment required to satisfy the complete set of grid constraints.

### 5.4.2 Case Study 2: Large Energy Systems

In this second case study, the MILP model is expanded to incorporate a wider range of components. The portfolio from the first case study, consisting of a PV system, BESS, and an electrical grid connection, is augmented with a CHP plant and a PEM electrolyzer. The heat and natural gas produced by these components can be sold to the grid.

Upon performing the initial, unconstrained DER optimization, the resulting line load time series is qualitatively similar to that observed in Case Study 1. However, a notable distinction emerges following the first application of grid constraints. In stark contrast to the previous case, a single iteration is sufficient to bring the line utilization below the 85 % limit for nearly all time steps, as depicted in Figure 5.4.

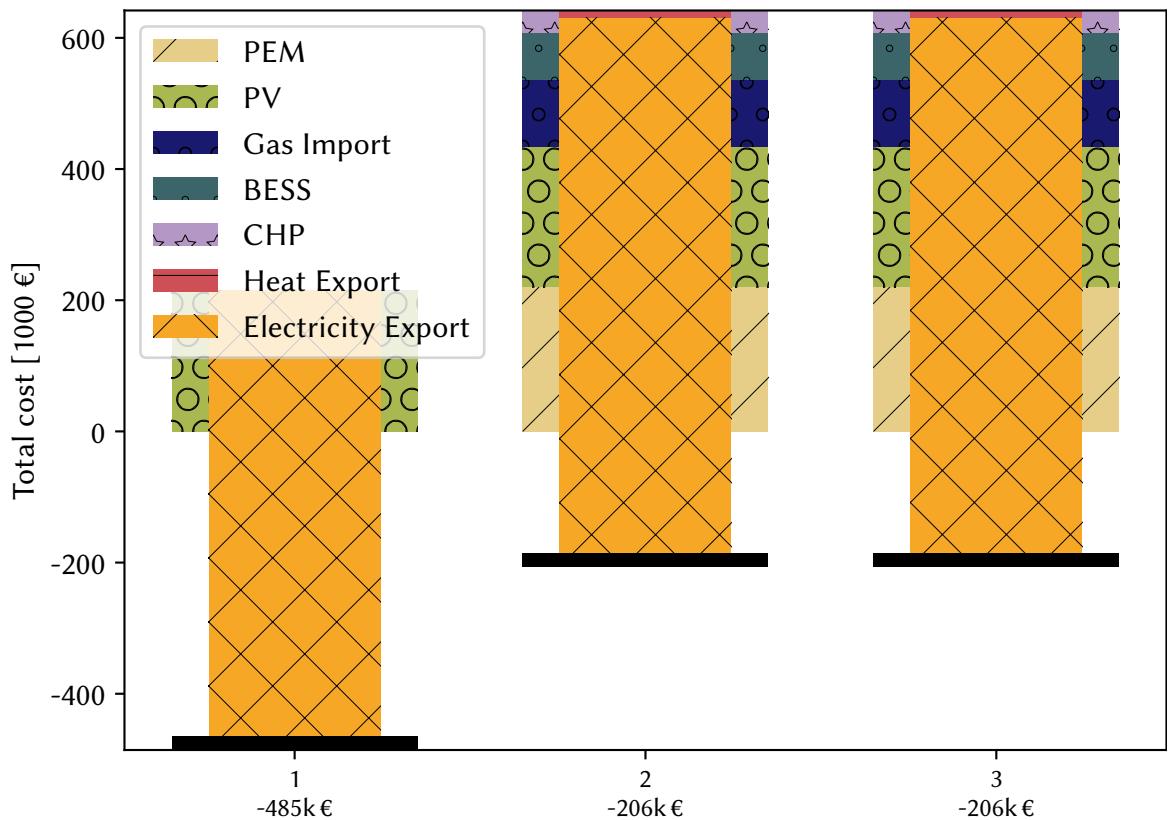


**Figure 5.4:** Line load at the start, after the first, and after the second iteration.

In the second iteration, only minimal adjustments to the constraints are necessary, as the DER now almost complies with the 85 % utilization limit. The algorithm terminates after this iteration as the line utilization is below 85 % at every time step. The final composition of the energy system is presented in Figure 5.5.

In the initial, unconstrained step, the optimal DER configuration requires no battery, PEM, or CHP plant. The total cost of operation for one year, including asset depreciation, is -485k €, which means that the DER is profitable and comparable to the initial state of Case Study 1. Afterwards, the algorithm makes a significant adjustment to meet the new grid constraints, incorporating all available components into the DERs design. This increases the total cost to -206k €, reducing the system's profitability but maintaining a net profit. In the final iteration, the new constraints are only marginally stricter. The DER configuration from the first iteration possesses sufficient flexibility to meet these new demands without requiring changes to its composition or incurring additional costs. The total cost, therefore, remains at -206k €.

A comparative analysis reveals the superior economic performance of the DER design from the second case study. This enhanced cost-efficiency is primarily attributable to the substantial capital cost of the large BESS necessitated by the grid constraints in Case Study 1. In contrast, the expanded component portfolio in the second case study provides more economic options to provide flexibility. The Power-to-Gas (PtG) plant serves as a dispatchable load, which not only aids in grid support but also generates additional revenue through the sale of the produced methane to the gas grid. Furthermore, the CHP plant provides on-demand power generation. The synergistic operation of these components substantially diminishes the required BESS capacity, as visually confirmed in Figure 5.5.



**Figure 5.5:** Composition of the energy system after each optimization step. Costs increase as constraints increase. This figure is adapted from [Gra+24a].

## 5.5 Discussion

To validate the efficacy of the proposed method, two case studies are conducted. This allows for an assessment of the algorithm's performance across different hardware configurations. The developed method is flexible and can be applied to various grids and DER compositions. However, the most relevant use-case for multi-energy DERs is within 20 kV medium-voltage distribution grids, close to or at the 400 V feeders. This focus is motivated by the growing trend of renewable generation in the 400 V grids and demanding EV charging infrastructure, directly integrated into 20 kV distribution grids. The successful application in these scenarios serves to demonstrate the algorithm's practical viability and immediate relevance.

In both case studies, the initial configuration of the DER is optimized exclusively for economic objectives, maximizing revenue at the expense of fulfilling requirements from the local distribution grid. A BESS is omitted from this configuration because the difference between peak and off-peak electricity prices makes price arbitrage infeasible, and the storage system itself lacks an independent revenue stream. In the subsequent design iteration, a minimally sized BESS is incorporated into the design to meet grid constraints. However, in scenarios where grid support demands exceed the power or energy rating of the small BESS, the optimization algorithm must dispatch additional assets. These secondary assets include a CHP plant and a PEM electrolyzer.

Although this process yields thermal energy (heat) from the CHP and hydrogen from the PEM electrolyzer, the revenue generated by selling these products is limited. Furthermore, the cost of buying methane from the gas grid substantially outweighs the income from selling hydrogen to the grid. Therefore, dispatching the CHP and PEM to ensure grid compliance incurs significant operational costs, highlighting a direct conflict between providing grid support and operating economically. By enabling the integration of new renewable generation or high-demand consumers, the additional investment in a "grid-friendly" DER may be justifiable due to substantial capital savings realized from obviating the need for expensive grid infrastructure upgrades or restrictions on the connected appliances.

While the presented method calculates the economically optimal DER composition that satisfies the imposed grid requirements, some limitations need to be acknowledged. As the calculation of the target load curve of the DER in Algorithm 3 is based on the line utilization, it is only driven by active power.

The computational runtime of the algorithm is variable, as it is dependent on three primary factors: the specific problem instance, the number of iterations required for convergence, and the underlying complexity of the MILP model used to represent the DER. For the case studies, each iteration required approximately two minutes to complete on a standard office computer equipped with an AMD 4750G processor.

As the performance of this algorithm is not fully optimized, significant reductions in computation time should be achievable with further algorithmic refinements. A key area for enhancement lies in the initialization of Algorithm 3. Instead of a naive starting point, employing an estimation technique to approximate the maximum utilization would provide a more informed initial guess. This would substantially accelerate the convergence of the search procedure, thereby reducing overall computation times.

## 5.6 Conclusion and Outlook

This chapter introduces a novel methodology for optimizing the design of DERs to include operational constraints derived from the local distribution grid. The core of this contribution is an iterative algorithm that combines the MILP-based DER design with power flow analysis. This process generates grid-dependent constraints that guide the optimization toward a multi-energy DER configuration capable of mitigating line congestions. The iterative nature of the algorithm is critical. It ensures that operational changes in the DER intended to solve a constraint violation do not inadvertently create new violations at other times. Through two case studies on an exemplary distribution grid, it is demonstrated that the proposed algorithm successfully computes a DER design that maintains line utilization below a predefined threshold. This grid-supportive capability comes at an increase in total cost, which converges as the algorithm iteratively refines the design until all grid constraints are satisfied.

Future work should focus on three key areas. First off, the algorithm should be extended to incorporate a broader range of grid parameters, such as voltage limits and transformer loading. Secondly, the algorithm's computational efficiency should be enhanced to reduce the computation time. Finally, the impact of the placement of the DER within the grid on the overall cost and efficacy of the solution should be investigated.



The run time of the algorithm discussed in the previous chapter is a major barrier to its practical application. The execution time of the algorithm is mainly driven by the slow, numerous power flow calculations. While machine learning techniques are commonly employed as surrogates for power flow calculations, this chapter proposes a more holistic, end-to-end neural network architecture. Two distinct GNN models are developed and evaluated: a two-stage, pre-trained model utilizing power flow calculations as a first stage and a single-stage, cold-start model. This analysis addresses the fifth research question: **How effective is a pre-trained GNN compared to a cold-start GNN for inferring the optimal DER power?**

The research presented in this chapter builds substantially upon the publication “Augmented Pre-trained Graph Neural Networks for Grid-Supportive Flexibility Control” [Gra+25]. While the core methodology is retained, this chapter provides a significantly extended and more comprehensive evaluation of the proposed approach.

## 6.1 Introduction

Predictive scheduling of flexible resources and fast decision-making processes are essential for the reliable operation of modern distribution grids. These require fast power flow calculations [De +22]. Power flow calculations enable the identification and mitigation of congestion scenarios through the dispatch of flexible loads or generators based on the line utilization, nodal voltage magnitude, and phase angle. Section 2.4 offers a comprehensive overview of power flow analysis.

An alternative approach to solving power flow calculations is the employment of machine learning techniques [PR02]. Neural networks can be trained to approximate the input-output mapping of a conventional power flow solver, such as Newton-Raphson. Once trained, the neural network can approximate power flow solutions at a fraction of the computational cost of the original solver. Contemporary research further advances this approach through the application of GNNs [De +22; Lia+22; Lin+24]. A defining characteristic of these graph-based neural networks is their ability to learn the intrinsic topological properties and physical foundations of power grids. This capability enables them to generalize and accurately infer power flow solutions even for grid topologies not included in the training dataset.

While the iterative algorithm presented in the previous chapter is effective, it is computationally expensive. At the same time, it is based on power flow calculations. Recent literature

demonstrates that GNNs can approximate power flows with high accuracy. This motivates researching a GNN that can be trained to function as a surrogate, capable of accurately approximating the results of the optimization algorithm at much lower computational cost. The present thesis presents a novel GNN architecture that extends beyond conventional power flow approximation. The GNN not only infers the power flow solution but also directly infers the active power required from a DER to mitigate an electricity line congestion. This latter, downstream task is accomplished by augmenting the power flow approximating GNN with more neural network layers that approximate the load or generation needed from a DER. The downstream task is one of many grid optimization and planning tasks based on power flow calculations. Other possible tasks include the calculation of the available load shifting capability of a grid, the minimization of power losses in the grid, and the optimization of the placement of DERs [SS17].

The central hypothesis investigated in this chapter is that employing a pre-trained GNN specialized in power flow inference as a first-stage model will both accelerate the training process for the downstream model and enhance predictive performance compared to a monolithic model trained end-to-end from scratch. To achieve robust generalization, the models are trained on heterogeneous dataset configurations. The individual datasets, which comprise multiple electrical grid topologies and realistic load scenarios, are combined in configurations designed to test the model's predictive accuracy on previously unseen network configurations. While pre-trained GNNs are an established technique [Liu+24], their application in the energy domain is a novel contribution of this work. Furthermore, this work uses realistic data. Unlike prior studies that often rely on synthetic profiles derived from statistical load distributions [BRZ19; De +22; Lin+24], the present work utilizes authentic SimBench load profiles and grid topologies as well as real-world grids. This allows for a realistic performance evaluation of the developed architectures.

This chapter continues as follows: The subsequent section outlines work related to previous applications of Artificial Neural Networks (ANNs) in power systems and discusses research on two-stage training of machine learning models. Afterwards, the developed GNN design and the process for generating training data are presented. This includes the design of a single-stage baseline model for comparison. In the evaluation, this baseline model is compared to the pre-trained model for three distinct dataset configurations. The results of this evaluation are then discussed in the final section of this chapter, answering the final research question: **How effective is a pre-trained GNN compared to a cold-start GNN for inferring the optimal DER power?**

## 6.2 Related Work

The application of ANNs to power system analysis dates back over three decades [Kal00; SP89]. Early methods employ ANNs to solve power flows for specific grid topologies. A representative approach involves training a neural network using the diagonal elements of a grid's admittance matrix as input features to predict the power flow solution [PR02]. A significant limitation of this paradigm, however, is its inherent lack of transferability. A model trained in this manner is limited to a single grid topology, and its generalization capability is restricted to unknown load scenarios within that topology.

GNNs have recently found application in the field of power system analysis [De +22; Lin+24]. This class of neural networks is specifically designed to leverage the topology of graph-structured data, enabling them to learn underlying physical properties while being invariant to the permutation of nodes and edges [Sca+09]. To approximate power flow solutions while explicitly exploiting the spatial nature of the grid, spatial-based GCNs have become a widely adopted approach [BRZ19; De +22; Lia+22]. The operational principle of these GNNs is a message-passing algorithm that facilitates the diffusion of information across the graph. In this process, which unfolds over multiple iterations corresponding to neural network layers, each node iteratively exchanges and aggregates information with its neighbors up to a distance of  $K$  hops. The suitability of this architecture for power system analysis is self-evident, as power grids can be intuitively represented as graphs where generators and loads constitute the nodes, and transmission lines form the edges, see Section 2.5.

Conventional power flow and optimal power flow solvers often rely on iterative optimization algorithms [De +22], which impose a significant computational burden, particularly for large power grids. To address this challenge, the work by de Jongh et al. [De +22] utilizes a physics-informed GNN to estimate grid states with high accuracy. The authors specifically highlight an architecture comprising TAGConv [Du+17] layers, noting its ability to achieve low MSE values even under conditions of incomplete sensor data.

In a recent publication, Lin et al. introduce the PowerFlowNet GNN architecture for the approximation of power flows [Lin+24]. Their method achieves substantial computational speedups for AC power flow calculations while delivering accuracy superior to that of conventional DC power flow approximations. The model contains five layers with a hidden dimension of  $d = 512$ , and utilizes a message-passing scheme that propagates information over three hops ( $K = 3$ ). The edge features, consisting of the line resistance and reactance, are aggregated by the one-hop message passing step. Information about voltage magnitude, voltage angle, injected active power, and reactive power is then collected from the whole neighborhood using the K-hop TAGConv. Lin et al. state that by combining the two principles in the PowerFlowNet layer, it is possible to approximate power flows with high quality [Lin+24]. Their evaluation includes three grid topologies and four power flow approximation techniques, next to the PowerFlowNet model. While their developed

architecture has longer run times than other machine learning architectures, such as a three-layer MLP or a three-layer GCN, the errors are consistently lower.

Two-stage training methodologies are a well-established approach in machine learning [Hus+18; Sha+21; Tzi+20]. The initial stage involves pre-training a portion of a neural network to capture foundational data representations or behaviors. A common technique for this pre-training phase is imitation learning [Hus+18]. In this paradigm, an agent is trained on instances of expert actions to learn a policy for solving a given task, as detailed in [Hus+18].

In the second training stage, the GNN is augmented with additional layers, while the weights of the pre-trained layers are frozen. These newly appended layers are then trained to perform a more specific downstream task by building upon the foundational features learned in the first stage. The objective is to refine the learned representations for a specialized application or to model more complex relationships that are not included in the initial pre-training. This two-stage approach is effective because the initial layers provide a pre-processing, simplifying the problem for the subsequent layers, which then only need to learn a mapping from these high-level features to the final output [Sch15].

The concept of pre-training a model on broad data for subsequent adaptation to a range of downstream tasks is conceptually aligned with the paradigm of foundation models, as defined by Bommasani et al. [Bom+22] and Hamann et al. [Ham+24]. However, it is crucial to note that the model presented herein diverges from this definition in two fundamental aspects. First, the presented model does not employ the self-supervised training methodology characteristic of foundation models. Second, it possesses a substantially lower number of trainable parameters. Consequently, while sharing a similar high-level philosophy, the present approach does not satisfy the key criteria necessary to be classified as a foundation model under the definition of Bommasani et al. [Bom+22].

## 6.3 Method

In the following, the architectures of the GNNs, including both a pre-trained model and a cold-start model for baseline comparison, are outlined. The method employs a two-stage process. First, a GNN is pre-trained to infer power flow solutions across various grid states. Subsequently, this base model is augmented with additional layers and fine-tuned for a specific downstream task: approximating the active power generated or drawn by a DER needed to mitigate line congestion.

### 6.3.1 Training Data

To ensure the models are trained on a diverse and realistic set of conditions, the training data is drawn from both the SimBench benchmark datasets grids 1-MV-urban, 1-MV-comm, and 1-MV-rural [Mei+20], and a real-world Medium Voltage (MV) distribution grid from the city of Kiel. The SimBench dataset offers realistic grid topologies and full-year load and generation profiles at a 15-minute temporal resolution. Each time step from these profiles is treated as an independent sample for model training.

The generation of the training dataset involves a two-part procedure. Initially, a power flow calculation is performed for every time step using the Newton-Raphson solver as implemented in Pandapower to establish the ground truth grid state. In any instance where this analysis reveals line utilization exceeding a predefined limit, a secondary computation is executed. The iterative procedure detailed in Algorithm 2 is then employed to determine the precise active power required to mitigate this congestion.

This process yields a comprehensive set of features for each node and edge, structured as follows:

- **Node Features:** For each node, a feature vector is constructed, comprising its voltage magnitude (p.u.), phase angle, active power, and reactive power. To support the downstream task, this vector is augmented with two additional features: the target mitigation power as calculated by Algorithm 2, and a one-hot encoding that identifies the specific node where the DER is located.
- **Edge Features:** Similarly, for each edge, the feature set includes its impedance and a one-hot encoding indicating which line is the target for reducing the utilization.

In accordance with published best practices for machine learning model development, data preprocessing is implemented [Jun22]. This includes the normalization of relevant input features. The process consists of two sequential steps. First, the data is centered by removing the sample mean. For each feature vector  $x$ , its mean  $\hat{x}$  is calculated across all samples and subsequently subtracted from every observation within that feature, as described by Equation (6.1). This operation ensures that each feature has a resulting mean of zero.

$$\begin{aligned}\hat{x} &= \frac{1}{n} \sum_1^n x^{(n)} \\ x^{(n)} &= x^{(n)} - \hat{x}\end{aligned}\tag{6.1}$$

Second, the features are scaled to have unit variance. Each feature value is divided by the standard deviation  $\sigma$  of that feature. The standardization for a given feature vector  $x$  can be expressed as

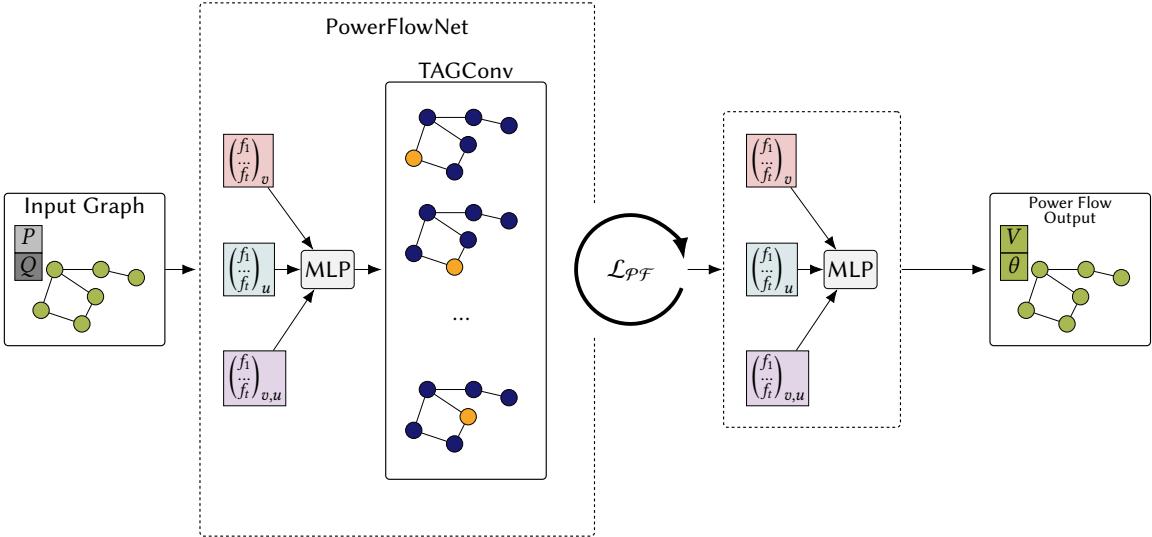
$$\begin{aligned}\sigma(x)^2 &= \frac{1}{n} \sum_1^n (x^{(n)})^2, \\ x^{(n)} &= \frac{x^{(n)}}{\sigma(x)}.\end{aligned}\tag{6.2}$$

These node and edge features are then compiled into arrays, forming the final input data for the model.

### 6.3.2 Power Flow Model

The model's architecture is centered around a stack of  $\mathcal{L}_{\mathcal{P}\mathcal{F}}$  PowerFlowNet layers [Lin+24]. Each of these layers integrates a two-layer MLP for message processing with a TAGConv operator for graph convolution [Du+17]. The MLP aggregates features of the immediate neighbors and adjacent edges. A comprehensive overview of the TAGConv layer can be found in Section 2.5.1. The initial layer processes an input feature vector of dimension two, corresponding to the active ( $P$ ) and reactive ( $Q$ ) power at each node, and maps it to a higher-dimensional representation of size  $d$ . The hidden dimension  $d$  is a hyperparameter that determines the model's number of learnable parameters, which in turn dictates the trade-off between computational performance and inference accuracy. A larger  $d$  increases the model's memory footprint and inference latency. Furthermore, an insufficient  $d$  leads to underfitting, where the model cannot accurately represent the relationships within the datasets. On the other hand, an excessive  $d$  leads to overfitting, where the model memorizes training dataset artifacts, resulting in poor performance on the validation set. After the last TAGConv layer, a two-layer MLP reduces the dimension to the final output dimension of two.

To determine the optimal model configuration, a comprehensive hyperparameter sweep is conducted. This sweep explores variations in the number of layers ( $\mathcal{L}_{\mathcal{P}\mathcal{F}}$ ), the hidden dimension ( $d$ ), and the size of the neighborhood ( $K$ -hop number). The training procedure is run for 100 epochs with a learning rate of  $\alpha = 0.001$ . This learning rate is recommended in the original publication of the Adam optimizer [KB17] as a “good default setting”. Fur-



**Figure 6.1:** GNN architecture of the power flow model. A two-layer MLP aggregates the features of the adjacent nodes and edges. The TAGConv incorporates all features of the nodes within the distance  $K$ . The voltage magnitude  $|V|$  and the voltage angle  $\theta$  are approximated for every node in the grid.

thermore, the PowerFlowNet [Lin+24] also implements this learning rate. To mitigate overfitting, an early stopping mechanism is employed. Throughout the training process, the model's performance is continuously evaluated on the validation dataset. The training epoch where the model achieves the best performance (e.g., lowest validation loss) is saved for the final evaluation. Consequently, the model utilized for testing is the one with the highest validated generalization performance, which does not necessarily correspond to the model from the final training epoch. An overview of the GNN architecture for power flow inference is depicted in Figure 6.1.

To facilitate comprehensive hyperparameter testing, several datasets are aggregated into distinct configurations. Each of the individual datasets contains approximately 35,000 time steps, corresponding to a full year of operational data recorded at a fifteen-minute resolution. To facilitate a realistic evaluation, a dataset from Kiel, a city in northern Germany, is also included. The dataset stems from the local DSO in Kiel and is not available publicly.

The following three distinct dataset configurations **C1**, **C2**, and **C3** are formulated:

- C1** The training and validation data contain the 1-MV-comm and 1-MV-urban SimBench grids with approximately 70,000 individual time steps. The final performance evaluation (test dataset) is conducted using the real-world dataset. The objective of this configuration is to evaluate the generalization performance of the models, specifi-

call their ability to transfer knowledge from the synthetic SimBench datasets to the empirical, real-world data from Kiel.

- [C2] Training and validation are conducted using the 1-MV-comm, 1-MV-urban, and 1-MV-rural SimBench datasets and the real-world dataset from Kiel with a total of about 735,000 data points. The test dataset consists of the 1-MV-comm grid and the real-world data. This configuration serves as a benchmark and is expected to yield the highest model accuracy because it contains the largest training dataset and a test set comprised exclusively of known grid topologies.
- [C3] In this configuration, the models are trained and validated exclusively on data from the real-world dataset. Subsequently, performance is evaluated on a test dataset composed solely of the 1-MV-comm and 1-MV-urban grids from the synthetic SimBench dataset. The objective of this configuration is to assess the model's ability to generalize from empirical data to a synthetic environment. This serves as a complementary analysis to the first configuration, which evaluates the sim-to-real transfer.

For all dataset configurations, the data is partitioned into training and validation sets using a 70/30 split, allocating 70 % of the data for training and the remaining 30 % for validation. The following Table 6.1 sums up all three configurations with their respective datasets.

Configuration	Training / Validation	Test
C1	1-MV-comm 1-MV-urban	Kiel
C2	1-MV-comm 1-MV-urban 1-MV-rural Kiel	1-MV-comm Kiel
C3	Kiel	1-MV-comm 1-MV-urban

**Table 6.1:** Overview of the three dataset configurations for training, validation, and testing.

The results of the hyperparameter evaluation are summarized in the Table 8.1, 8.2, and 8.3 that can be found in the appendix. For clarity, each of these tables displays only the ten models that achieve the lowest test loss. In Table 6.2, the test and validation losses for each model are summed, and the table presents the ten models with the lowest aggregate loss.

$\mathcal{L}_{\mathcal{P}\mathcal{F}}$	$d$	$K$	Avrg. Validation MSE	Avrg. Test MSE	Avrg. Rank
3	256	3	1.24	1.67	7
2	512	4	0.22	3.87	9
5	128	2	1.09	2.81	9
4	256	5	1.19	2.17	10
3	192	2	0.18	3.10	13
4	256	4	0.46	8.25	13
4	192	3	0.81	259.54	13
3	128	2	0.52	1.91	14
2	256	4	1.25	2.79	14
3	512	5	0.11	20.48	14

**Table 6.2:** Performance average of various power flow model configurations on all three dataset configurations. The models are sorted by their average ranking across the tests on the three dataset configurations.

The evaluation phase highlights a trade-off between performance on the validation and test datasets. Different models excel in different metrics. Specifically, when examining the inference accuracy on the test set, the model with  $\mathcal{L}_{\mathcal{P}\mathcal{F}} = 3$  layers, a hidden dimension of  $d = 256$ , and  $K = 3$  hops demonstrates the lowest losses on average. In contrast, the model with  $\mathcal{L}_{\mathcal{P}\mathcal{F}} = 4$  layers, a hidden dimension of  $d = 128$ , and  $K = 2$  hops yields the lowest test loss average, while the model with  $\mathcal{L}_{\mathcal{P}\mathcal{F}} = 5$  layers, a hidden dimension of  $d = 192$ , and  $K = 3$  hops performs the best on the validation set. In selecting the ideal architecture for the first stage of the pre-trained model, a holistic performance perspective is critical. Therefore, the model with the highest average ranking in the hyperparameter tests on the three dataset configurations is chosen, as this indicator is interpreted as a sign of greater robustness across diverse dataset configurations and grid topologies, making it the most reliable candidate.

### 6.3.3 Model Architecture for Mitigating Congestion

To approximate the power required for congestion mitigation, two model architectures are developed and compared. One is pre-trained on power flow data, and a baseline model is trained from a cold start. The hyperparameter optimization for both models explores a comprehensive search space. This includes varying the hidden dimension  $d$  from 64 to 512, the number of hops ( $K$ ) from two to six, and the number of layers in the downstream task ( $\mathcal{L}_{\mathcal{D}\mathcal{E}\mathcal{R}}$ ) from two to seven.

The input data for the second stage of the pre-trained architecture necessitates the inference of the power flow model. The second training process, therefore, does not use ground-truth power flow data. Rather, it is based on the power flow approximation gener-

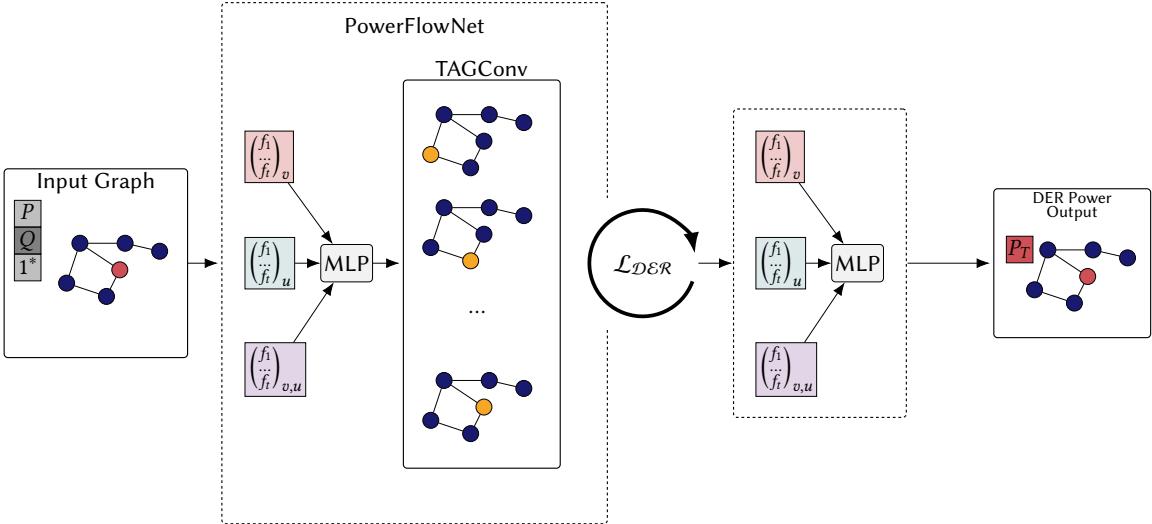
ated by the first-stage GNN. This implies that any approximation errors from the initial stage are propagated into the subsequent one. However, this is the case during training as well as during testing, which could potentially reduce the impact of the approximation error. Nevertheless, the implications of this inaccuracy are challenging to estimate in advance and are evaluated separately.

All hyperparameter tests and evaluations are conducted using the same power flow model with  $\mathcal{L}_{\mathcal{P}\mathcal{F}} = 3$  layers, a hidden dimension of  $d = 256$ , and  $K = 3$  hops. This model resembles the best trade-off between the models tested in the previous section, as can be seen in Table 6.2. The downstream, second GNN for inferring the power required for congestion mitigation consists of  $\mathcal{L}_{\mathcal{D}\mathcal{E}\mathcal{R}} = 2$  layers, a hidden dimension of  $d = 256$ , and  $K = 2$  hops, as this configuration achieves the highest average rank during testing (Table 6.3).

$\mathcal{L}_{\mathcal{D}\mathcal{E}\mathcal{R}}$	$d$	$K$	Avrg. Validation MAPE	Avrg. Test MAPE	Avrg. Rank
2	256	2	0.02	0.55	2
3	256	2	0.03	0.60	3
3	128	2	0.05	1.01	4
2	128	2	0.09	1.07	5
4	128	3	0.03	1.46	6
3	512	4	0.03	1.54	7
3	128	5	0.03	0.70	8
3	64	6	0.03	1.66	8
2	64	6	0.03	1.72	9
4	512	4	0.03	1.88	11

**Table 6.3:** Performance average of various pre-trained model configurations on all three datasets. The models are sorted by their average rank.

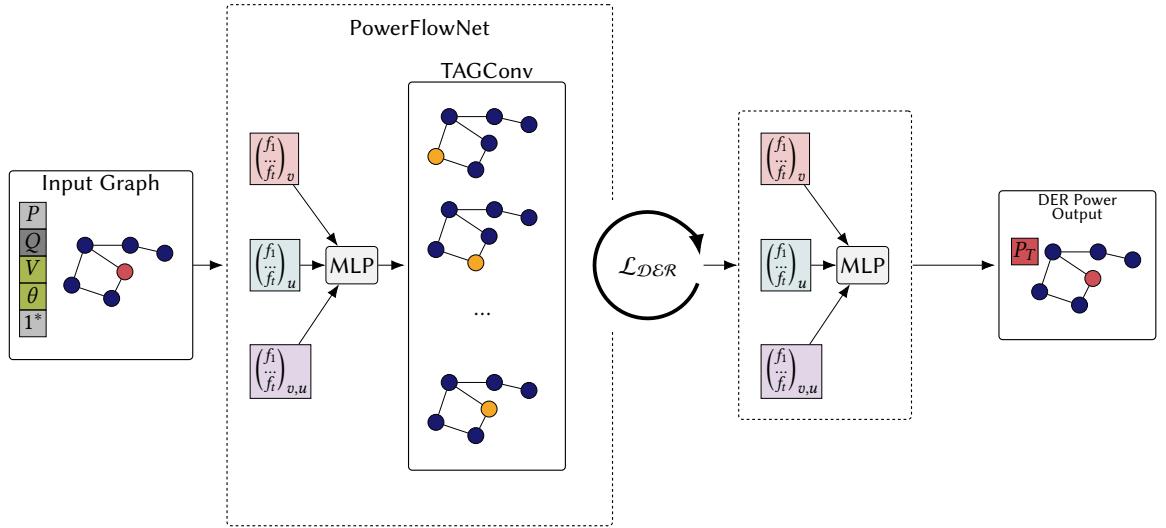
To optimize the training process, the datasets are filtered to include only time steps with utilization limit violations. This targeted sampling avoids heavily underestimating the needed DER power adjustment caused by the numerous samples where the target DER adjustment is zero, which could otherwise impede effective learning of the required non-zero power values. The training of the models is performed over 100 epochs, using the Adam optimizer with a learning rate set to  $l_r = 0.001$ . This number of epochs allows both models to converge to a minimum on their respective training losses. An extensive explanation of the Adam optimizer is given in Section 2.5.3. As with the power flow training, the model’s performance is continuously validated during training. The model of the epoch exhibiting the lowest validation loss is saved for evaluation on the test dataset. This early stopping prevents overfitting the model to the training dataset and results in a model with “hopefully better test set error” [GBC16].



**Figure 6.2:** GNN architecture of the pre-trained model for congestion mitigation. The green features of the nodes in the input graph denote features approximated by the power flow GNN, and the feature  $1^*$  denotes the one-hot encoded location of the DER. The MLP after the last TAGConv layer reduces the feature dimension to one. Although this feature is present at all nodes, only the one-hot encoded target node, colored in red, is of relevance.

The pre-trained model consists of a base power flow model and a downstream GNN, which acts as a surrogate for the iterative optimization algorithm. After the initial training of the power flow model, the model’s weights are fixed. The power flow model is depicted in Figure 6.1 and outputs voltage magnitudes  $|V|$  and phase angles  $\theta$  for all nodes in the graph. The task-specific second stage then learns to map a combination of the power flow model’s output and the initial input features (active power  $P$ , reactive power  $Q$ , and a DER location indicator  $1^*$ ) to the final output  $P_T$ . This second stage is depicted in Figure 6.2. The two-stage design is contrasted with the cold-start model (Figure 6.3), which is a single, end-to-end GNN trained directly on the same task.

As indicated in Table 6.4, the cold-start model configured with three layers ( $\mathcal{L}_{DER} = 3$ ), a hidden dimension of  $d = 128$ , and  $K = 3$  hops consistently outperforms all other hyperparameter configurations, achieving the best performance on each of the three dataset configurations, ranking number one in each of them.



**Figure 6.3:** GNN architecture of the cold-start model for congestion mitigation. The target node where the DER is located is colored in red.

$\mathcal{L}_{DER}$	$d$	$K$	Avrg. Validation MAPE	Avrg. Test MAPE	Avrg. Rank
3	128	3	0.03	0.77	1
3	512	3	0.02	1.21	2
4	512	3	0.01	1.35	3
3	256	4	0.02	1.40	5
2	256	4	0.04	1.41	5
5	128	3	0.08	1.50	6
5	512	3	0.02	1.54	7
4	256	3	0.03	1.60	7
3	256	3	0.03	2.04	10
2	256	5	0.03	2.17	12

**Table 6.4:** Performance average of various cold-start model hyperparameter configurations on all three datasets. The models are sorted by their average rank.

## 6.4 Evaluation

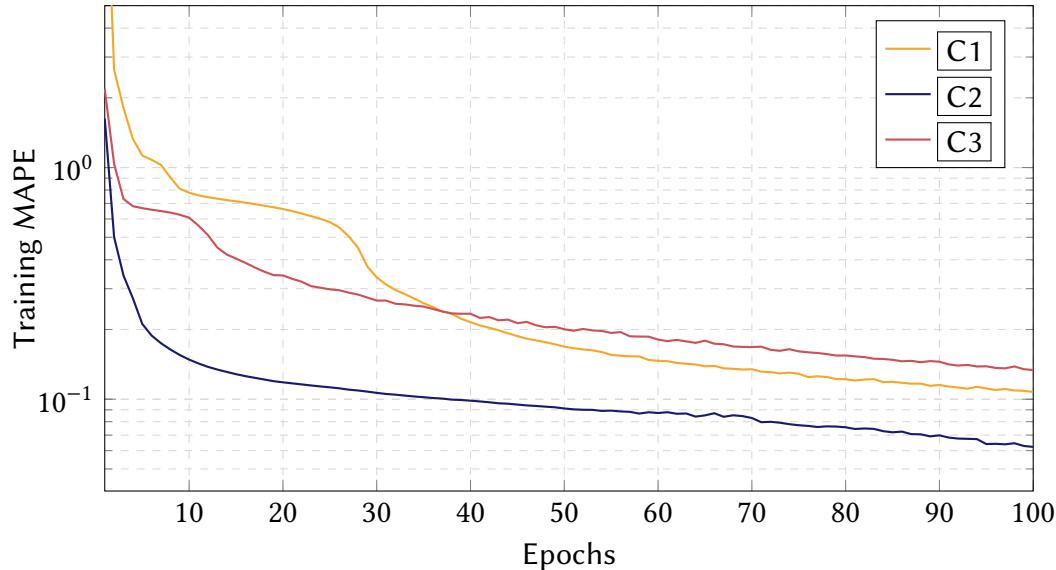
This section evaluates the best-performing model configurations identified during the hyperparameter search. The analysis focuses on the Mean Absolute Percentage Error (MAPE) on the test dataset, as it is designed to represent the real-world impact of the error.

### 6.4.1 Power Flow Model

During the hyperparameter search in the previous section, a three-layer model with a hidden dimension of  $d = 256$  and  $K = 3$  hops demonstrates sufficient performance across all evaluated datasets. Although it does not perform the best in every validation and test case, it performs well consistently across all dataset configurations, making it a broadly applicable solution.

Model	Test MAPE C1	Test MAPE C2	Test MAPE C3
$\mathcal{L}_{\mathcal{P}\mathcal{F}} = 3, d = 256, K = 3$	0.65	0.45	3.92

**Table 6.5:** Performance evaluation of the power flow inference model with the best performance during the hyperparameter tests on the three dataset configurations. All values are MAPE.



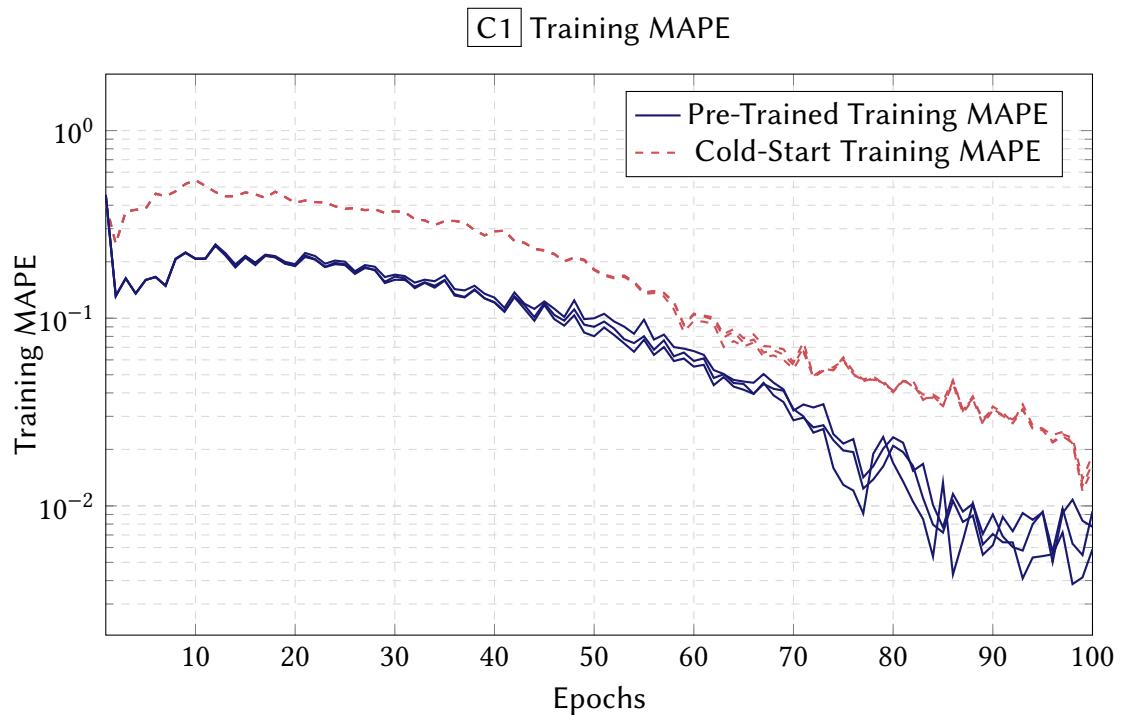
**Figure 6.4:** Training losses of the model with  $\mathcal{L}_{\mathcal{P}\mathcal{F}} = 3$  layers, a hidden dimension of  $d = 256$ , and  $K = 3$  hops.

The denormalized test MSE after 100 epochs of training this model ( $\mathcal{L}_{\mathcal{P}\mathcal{F}} = 3, d = 256, K = 3$ ) ranges from  $7.62 \cdot 10^{-3}$  p.u. for the first dataset configuration C1 to  $33.97 \cdot 10^{-3}$  p.u. for the third dataset configuration C3. As illustrated in Figure 6.4, the training score of the model on the third dataset configuration C3 is also the highest among the tested configurations. This finding is consistent with the observed behavior of other models during the hyperparameter search.

### 6.4.2 Training the Pre-Trained and Cold-Start Models

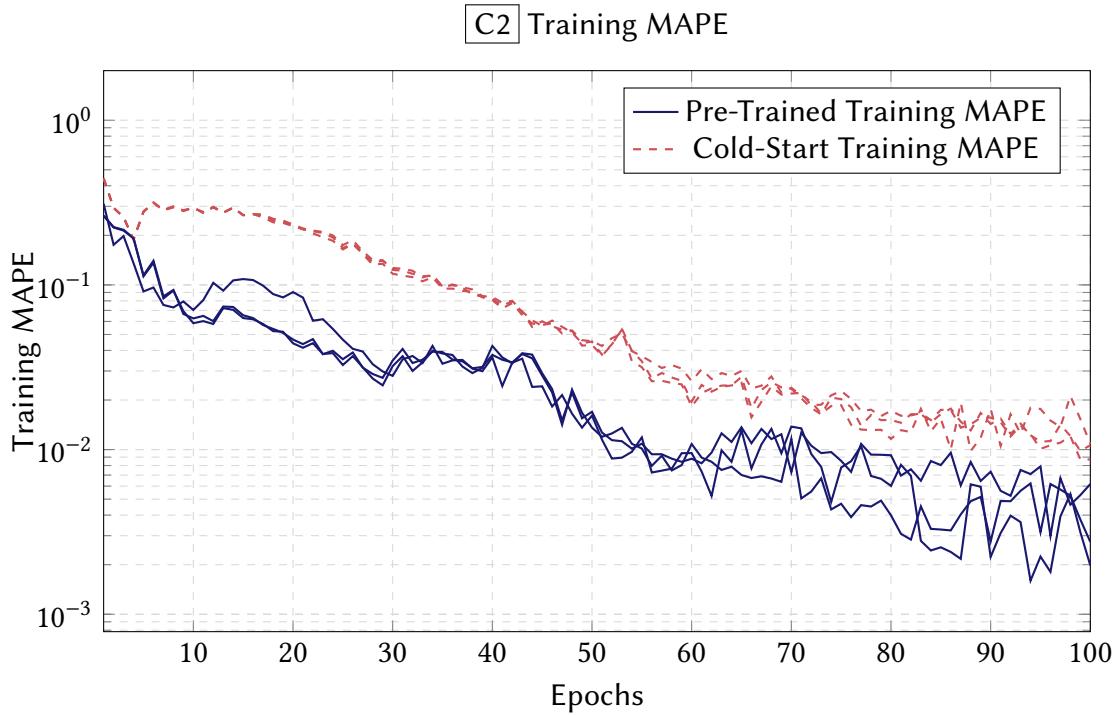
For both the pre-trained and the cold-start architectures, three training runs are performed on the same dataset configuration. Since all training procedures exhibit very similar behavior, no additional runs are conducted.

As shown in Figures 6.5, 6.6, and 6.7, the pre-trained model achieves better training losses than the cold-start model. This advantage is particularly evident in Figure 6.5, which shows that the MAPE of the pre-trained model trainings are lower than the MAPE of the cold-start model trainings after almost every epoch.



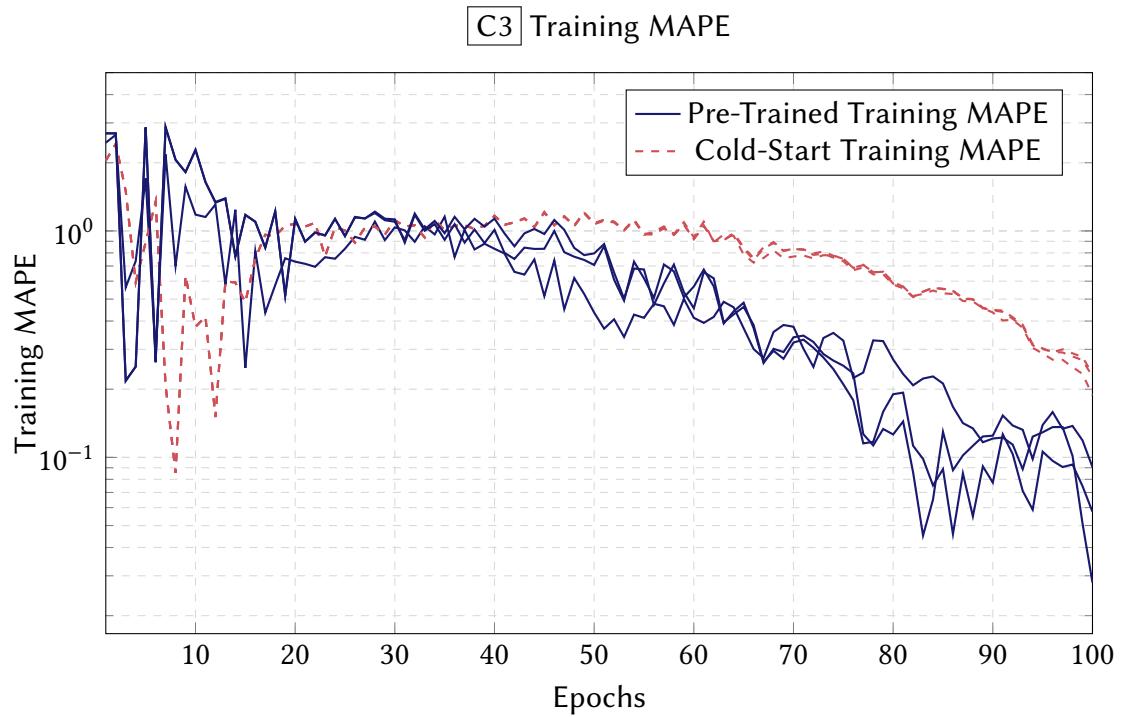
**Figure 6.5:** Comparison of the training losses on the first dataset configuration C1. The figure includes three exemplary runs from the pre-trained model and the cold-start model, respectively.

The second dataset configuration C2 generally shows high accuracies for both models. Between the individual runs of the same model, the loss is almost identical and the volatility is low, as can be seen in Figure 6.6.



**Figure 6.6:** Comparison of the training losses on the second dataset configuration C2 in three exemplary training runs.

The training process for the third dataset configuration C3 is characterized by significantly higher volatility compared to the other two. Figure 6.7 depicts the periodic fluctuations, during which performance temporarily declines and the loss increases. Still, a clear downward trend can be seen in the training and validation loss for both architectures. However, the pre-trained model remains superior with this configuration: after the entire training period of 100 epochs, it shows lower loss than the cold-start model.



**Figure 6.7:** Comparison of the training losses on the third dataset configuration C3 in three exemplary training runs.

#### 6.4.3 Performance Evaluation of the Pre-Trained and Cold-Start Models

A comparative analysis of the training, validation, and test losses of the pre-trained and cold-start models is conducted to evaluate their final performance. In addition, a third two-stage model is assessed as a benchmark. The model is supplied with accurate power flow calculations as input for its second stage, both during the training and the evaluation phases. This evaluation is designed to isolate the model's performance from any inaccuracies introduced by the first-stage power flow approximating GNN.

<span style="border: 1px solid black; padding: 2px;">C1</span> Model / Loss	Training (MAPE)	Validation (MAPE)	Test (MAPE)	MSE Test (W)
Pre-Trained	0.006 %	0.023 %	1.206 %	37.06
Cold-Start	0.024 %	0.045 %	1.454 %	49.97
Only Second Stage	0.005 %	0.017 %	0.906 %	20.93

**Table 6.6:** Performance evaluation on the first dataset configuration C1 after 100 epochs of training.

In the first data set configuration **C1**, the pre-trained model achieves better training, validation, and test scores than the cold-start model. As depicted in Figure 6.5, the pre-trained model has an advantage over the cold-start model throughout the whole training process. The pre-trained model can be further improved by utilizing a perfect power flow input. In this configuration, the use of a conventional power flow solver as the primary stage would result in a performance improvement.

<b>C2 Model / Loss</b>	<b>Training (MAPE)</b>	<b>Validation (MAPE)</b>	<b>Test (MAPE)</b>	<b>MSE Test (W)</b>
Pre-Trained	0.007 %	0.022 %	0.224 %	6.98
Cold-Start	0.010 %	0.020 %	0.21 %	48.39
Only Second Stage	0.004 %	0.004 %	0.058 %	0.51

**Table 6.7:** Performance evaluation on the second dataset configuration **C2** after 100 epochs of training.

For the second dataset configuration **C2**, the pre-trained model demonstrates superior performance in comparison to the cold-start model during the training process. Nevertheless, the final performance evaluation reveals what appears to be a contradiction in the test loss. The cold-start model shows a lower MAPE on validation and test data but a higher test MSE.

<b>C3 Model / Loss</b>	<b>Training (MAPE)</b>	<b>Validation (MAPE)</b>	<b>Test (MAPE)</b>	<b>MSE Test (W)</b>
Pre-Trained	0.011 %	0.009 %	0.233 %	157.79
Cold-Start	0.027 %	0.047 %	0.505 %	3,555.30
Only Second Stage	0.013 %	0.014 %	0.489 %	545.41

**Table 6.8:** Performance evaluation on the third dataset configuration **C3** after 100 epochs of training.

As shown in detail in Table 6.8, the evaluation of the third dataset configuration shows consistently higher MSE during testing compared to the other two dataset configurations. This increased difficulty is evident throughout the training process and results in a higher final test loss. This trend is particularly pronounced in the cold-start model, whose test MSE exceeds 3.5 kW. In this training scenario, the pre-trained model outperforms the cold-start model in all loss functions. This configuration leads to a counterintuitive result. The two-stage model using perfect power flow data as inputs for the second stage performs worse than the model that relies on estimates derived from the first-stage GNN.

## 6.5 Discussion

The power flow model achieves a test MSE between  $7.62 \cdot 10^{-3}$  p.u. and  $33.97 \cdot 10^{-3}$  p.u. Although these results are higher than those reported by Lin et al. [Lin+24], they are comparable to the results of De Jongh et al. [De +22]. This discrepancy is probably due to the type of data used: This study uses real data and realistic SimBench datasets, whereas Lin et al. train and test their model using synthetic data with a specific distribution [Lin+24]. Furthermore, the model architecture was chosen for its ability to generalize across different datasets, rather than being specialized for optimal performance on a single dataset. As shown in Table 6.5, such specialization would have led to better performance for individual dataset configurations, underscoring the deliberate trade-off between maximum accuracy for individual use-cases and good performance on all evaluated dataset configurations.

While the hyperparameter evaluation includes the specific GNN architecture proposed by Lin et al. [Lin+24], comprising four layers ( $\mathcal{L}_{\mathcal{P}\mathcal{F}} = 4$ ), a hidden dimension of  $d = 128$ , and  $K = 3$  hops, it does not yield the best average performance on the evaluated datasets. Alternative model configurations consistently achieve lower losses. This finding may be attributed to the increased complexity of the datasets; the various grid topologies and realistic load scenarios used in the present study are a more challenging environment than the one investigated by Lin et al..

Compared to the previously published work “Augmented Pre-trained Graph Neural Networks for Grid-Supportive Flexibility Control” [Gra+25], the present work incorporates a real-world dataset from the German city of Kiel. Consequently, further hyperparameter tests are conducted to identify the model that optimizes the trade-off between the three dataset configurations. The results of the experiments on the real-world dataset in this study demonstrate the models’ capability to infer power flows on real-world grids, thereby enhancing the method’s applicability.

A training duration of 100 epochs proves to be sufficient in all cases for the power flow, the pre-trained models, and the cold-start models. Although the MAPE on the training dataset continues to improve until the 100th epoch, further training leads to overfitting, which degrades the accuracy on the validation data and confirms that a limit of 100 epochs is appropriate.

The results for the first data set configuration C1 (Table 6.6) are in line with theoretical expectations. While the pre-trained model slightly outperforms the cold-start model, its performance is limited by the accuracy of the first-stage power flow GNN. Consequently, the two-stage benchmark model, which uses perfect power flow data in the first stage and is therefore not subject to power flow approximation inaccuracies, achieves the lowest test loss. This demonstrates the impact of the error propagated from the first stage.

For the second dataset configuration [C2], lower overall losses are expected due to the more extensive training dataset and a test dataset consisting exclusively of topologies contained in the training data. This hypothesis is confirmed by the final test results, which show a consistent reduction in losses across all models (Table 6.7). An exception is the MSE test loss for the cold-start model, which is comparable to the accuracy in the first dataset configuration [C1]. This discrepancy may be due to rare but significant errors that lead to a higher MSE loss, while at the same time, a comparatively low MAPE loss is observed. Interestingly, the performance gains of the pre-trained models are specific to the test phase, as both the training and validation losses are comparable to those of the first configuration. Only the cold-start model achieves significantly better losses during the training and validation phase than in the first data set configuration.

The evaluation of the results of the third dataset configuration [C3] shows significantly lower test accuracy for all models (Table 6.8). This drop in performance can be attributed to a fundamental challenge in generalization: models trained exclusively on a single real-world grid topology are unable to adapt to the SimBench grids. As this challenge only exists in the test phase, the training losses are comparable to those of the second data set configuration. The resulting discrepancy between training and test performance is therefore to be expected and underscores the difficulty of generalizing from such a limited training data set.

The iterative algorithm results range from 0 MW to 5 MW. The pre-trained model demonstrates adequate inference capability with a test MAPE of 1.206 % and test MSE of 37.06 W on dataset configuration [C1]. Performance improves significantly on dataset configuration [C2], achieving a test MAPE of 0.224 % (6.98 W). Although the MAPE on the test dataset of dataset configuration [C3] is at 0.233 % and in line with the results on the other dataset configurations, the MSE is at 157.79 W. These results demonstrate that learned representation transferability exhibits variable difficulty. Optimal training necessitates comprehensive datasets with diverse topological configurations to achieve maximum performance, as demonstrated by configuration [C2].

## 6.6 Conclusion and Outlook

This chapter investigates GNNs for approximating the minimum DER power (load or generation) required to relieve the grid, as formulated in Chapter 5. With a single-stage model that directly infers the optimal DER power and a two-stage architecture that first approximates the power flows and then, in a subsequent stage, approximates the optimal DER power based on these approximations, a comparison is performed to evaluate the advantages and limitations of these two approaches.

This evaluation focuses on the generalizability of the GNN models, necessary for their practical use in real-world power grids. The accuracy of the GNN models trained on a specific grid topology is assessed on different, previously unseen topologies. To this end, a GNN architecture for power flow approximation used as the first stage is benchmarked across diverse datasets using three distinct configurations. The evaluation contrasts with prior work by Lin et al. [Lin+24] in which a comprehensive evaluation of generalization capabilities across multiple grid topologies is missing. The evaluation conducted in the present work leverages both real-world operational data from an active distribution network and established synthetic SimBench grids to ensure a realistic assessment.

The evaluation shows that the power flow inferring model achieves similar accuracy compared to other models in the literature. Especially when trained on a large dataset and with tests that are only conducted on grid topologies that are also included in the training dataset (C2), the accuracy of the power flow inference is high with a MAPE of 0.45 % (Table 6.5).

Implementing the power flow inference model as an initial stage, the pre-trained variant consistently outperforms the cold-start approach in both training and testing across all dataset configurations. During training, the cold-start model converges more slowly, requiring additional epochs to match the inference accuracy of the pre-trained model. Consistent with the power flow inference model training, utilizing SimBench grids and real-world grid data yields the maximum dataset size while eliminating unseen topologies during testing, thereby achieving high approximation accuracy. Conversely, evaluation on novel grid topologies not present in the training data leads to diminished accuracy (C1 and C3).

The results demonstrate that both pre-trained and cold-start models can reliably infer the required DER power on real-world grids with improved performance when the corresponding grid topologies are included in the training dataset. However, the accuracy of the inference depends highly on the dataset used for training, underscoring the paramount importance of training data selection and the explicit evaluation of model generalizability.

In this chapter, it is demonstrated that a pre-trained GNN outperforms a cold-start GNN in inferring the optimal power required from a DER to reduce line utilization. In most of the evaluated test cases, the pre-trained model shows faster training times and more accurate inference results after 100 epochs of training. This finding answers the fifth research question: **How effective is a pre-trained GNN compared to a cold-start GNN for inferring the optimal DER power?**

Further research is required to explore additional use cases of pre-trained GNN in power grids. While this work highlights the potential and outlines possible applications, their implementation and evaluation remain open tasks. Moreover, assessing the method on a more comprehensive dataset that includes a larger set of real-world grids would provide stronger evidence of its practical applicability.



The proliferation of renewable electricity generation in the distribution grids and the electrification of the transport and heating sectors introduce decentralized generation and high-power loads. Legacy grid operation without control mechanisms and limited grid monitoring is akin to navigating a complex system blind. To ensure grid stability and economic asset management, a transition towards comprehensive grid monitoring is essential. A cost-effective alternative to utility-grade measurement devices emerges in increasingly popular smart home hardware, such as smart plugs.

To assess real-world applicability, a telecommunication backend is designed and implemented in Chapter 3. By employing standardized database servers and message brokers, secure communication between smart home devices and the backend is ensured. This infrastructure enables continuous collection of voltage measurements at the socket outlets where the smart plugs are connected. Analysis of the 10-second interval measurements reveals two error components: a constant offset and a varying error. Averaging measurements over extended periods and comparing them with a reference device allows estimation of the constant offset, while the varying error is shown to follow a normal distribution. Correcting the constant offset yields measurement accuracy of the smart plugs sufficient for grid monitoring, thereby addressing research question 1: **To what extent can smart plugs be used for distribution grid monitoring?**

Even inexpensive measurement devices, like smart plugs, are unlikely to be deployed at every node within the grid. This limitation underscores the need for pseudo-measurements to bridge the gaps in data collection. To be effective for grid monitoring, these measurements must be generated rapidly, be computationally inexpensive, and remain sufficiently accurate. The feature propagation algorithm presented in this thesis is able to generate pseudo-measurements. This motivates research question 2: **How viable is the use of a feature propagation algorithm for pseudo-measurement generation in distribution grids compared to other related methods?**

The feature propagation algorithm is evaluated against the GINN algorithm using an electricity grid graph. In scenarios with 50 % to 98 % of features missing, the feature propagation algorithm demonstrates superior performance. As the method of the feature propagation is closely related to heat propagation using the finite element method, it is not based on historical or synthetic data for training and is simple to compute.

The acquisition of comprehensive distribution grid monitoring data enables optimized planning and operation of DERs. A direct comparison of algorithms for optimizing DER schedules reveals a clear distinction between heuristic and analytical methods. When tested on a year-long schedule optimization, the analytical MIQP algorithm significantly outperforms heuristic approaches. The MILP and MIQP algorithms achieve a significantly lower RMSE and are approximately 500 times faster than the EA. However, this performance advantage introduces a trade-off in modeling effort, as the analytical MILP and MIQP methods require manual modeling of multi-energy DER components. In contrast, the EA utilizes component lookup tables generated from simulations. These findings collectively address the third research question: **What are the advantages and disadvantages of optimizing a multi-energy DER using an evolutionary algorithm, an MILP, and an MIQP?**

The application of the optimization algorithms can be further expanded to include the surrounding distribution grid area. Sizing the components of the DER to fulfill the needs of the DSO's congestion management often increases capital and operating costs of the DER, as it may require oversized components or suboptimal operation.

To resolve this compromise, Chapter 5 describes an algorithm that calculates the minimum power generation or load required by a DER to avoid grid congestion. Two case studies with different component portfolios evaluate the iterative algorithm. The algorithm maintains line utilization below 85 % at all time steps while identifying economically optimal component sizes. Total costs of the multi-energy DER increase drastically depending on specific case conditions. This holistic calculation answers the fourth research question: **How does the annual cost of a DER designed to provide grid supporting services compare to the annual cost of the most economical design?**

While the iterative planning process enables the optimization of a multi-energy DER, its practicality is hampered by the high computational expense. The power flow calculations can be substituted by a GNN surrogate model that infers results orders of magnitude faster than conventional numerical solvers. However, calculating the optimal power required to reduce line utilization below a certain threshold remains a significant computational bottleneck.

To eliminate the iterative calculation, a novel two-stage GNN architecture is developed in Chapter 6. This approach bypasses the iterative algorithm by directly approximating the power required to mitigate congestion. Its effectiveness is validated by comparing the two-stage, pre-trained model, which incorporates a power flow inferring model as its first stage, against a conventional single-stage, cold-start GNN. The simpler, single-stage model attempts to infer the required power directly without the pre-trained, multi-stage process. Evaluating the first-stage power flow model independently, the accuracy of the inference is found to approximate related publications. On previously unseen grid topologies, the MAPE is between 0.45 % and 3.92 %, corresponding to between  $7.62 \cdot 10^{-3}$  p.u. and  $33.97 \cdot 10^{-3}$  p.u.. Using the power flow model as the first stage, the pre-trained model outperforms the cold-start model across all dataset configurations. Averaged over the three dataset configurations, the pre-trained model achieves a test MAPE of approximately 0.55 %, compared to 0.72 % for the cold-start model. When a conventional power flow solver is employed as the first stage, the MAPE decreases further to about 0.48 %. The evaluation demonstrates that by splitting the complex inference task into a structured, two-stage process, the model can use the underlying power flow calculations to learn more effectively. It also results in more accurate and reliable predictions for real-time grid control and a significantly accelerated multi-energy DER sizing procedure. This answers the fifth research question: **How effective is a pre-trained GNN compared to a cold-start GNN for inferring the optimal DER power?** Furthermore, the performance of both models on real-world datasets demonstrates the practical feasibility of the proposed approach.

This thesis compiles and advances the results of the following publications:

- “Distribution Grid Monitoring Based on Widely Available Smart Plugs”, Grafenhorst et al., published as a conference paper at the Energy Informatics Academy conference in 2023 [GFH24b]
- “Poster Abstract: A three-phase electricity grid model of a single family house”, Grafenhorst et al., published as a poster at the DACH+ Conference on Energy Informatics in 2023 [GFH23]
- “Distribution grid monitoring based on feature propagation using smart plugs”, Grafenhorst et al., published as a journal paper in the Energy Informatics Journal in 2024 [GFH24a]

- “Heuristic vs. Analytical Energy Hub Optimization: Design, Implementation, and Trade-Offs”, Grafenhorst et al., published as a conference paper at the IEEE SEGE conference in 2024 [Gra+24b]
- “Grid Aware Portfolio Optimization of a Multi-Energy DER”, Grafenhorst et al., published as a conference paper at the IEEE iSPEC conference in 2024 [Gra+24a]
- “Augmented Pre-trained Graph Neural Networks for Grid-Supportive Flexibility Control”, Grafenhorst et al., published as a conference paper at the ACM E-Energy conference in 2025 [Gra+25]

While this thesis establishes a proof of concept, further investigation is necessary to validate the broader applicability and scalability of these techniques. Consequently, future work focuses on two aspects. First, an evaluation on a wider range of more complex and topologically diverse real-world grids is required. Assessing the performance and robustness of the GNNs on such grids is essential to confirm their practical viability. Second, the current analysis needs to be expanded to explore scenarios with multiple DERs located within the same distribution grid. Multiple DERs potentially allow for synergistic and conflicting interactions, requiring coordinated control strategies to manage multiple DERs holistically.

Rank	$\mathcal{L}_{\mathcal{P}\mathcal{F}}$	$d$	$K$	Validation MSE	Test MSE
1	3	256	4	0.10	0.47
2	3	512	5	0.06	0.52
3	5	128	2	0.12	0.54
4	2	512	4	0.18	0.55
5	5	256	3	0.03	0.55
6	3	128	2	0.15	0.59
7	2	256	4	0.19	0.61
8	3	256	3	0.11	0.65
9	4	192	3	0.06	1.33
10	4	128	3	0.10	1.34

**Table 8.1:** Performance of various power flow models using dataset configuration C1.

Rank	$\mathcal{L}_{\mathcal{P}\mathcal{F}}$	$d$	$K$	Validation Loss	Test Loss
1	4	256	4	0.02	0.33
2	5	256	4	0.03	0.38
3	4	192	3	0.03	0.41
4	4	256	3	0.02	0.43
5	3	256	3	0.03	0.45
6	5	256	5	0.03	0.48
7	2	512	4	0.17	0.49
8	4	128	4	0.03	0.54
9	4	256	5	0.02	0.58
10	5	256	2	0.02	0.58

**Table 8.2:** Performance of various power flow models using dataset configuration C2.

Rank	$\mathcal{L}_{\mathcal{P}\mathcal{F}}$	$d$	$K$	Validation Loss	Test Loss
1	4	128	2	4.63	0.75
2	4	256	2	0.26	2.10
3	3	128	2	1.35	2.79
4	3	192	2	0.34	2.84
5	3	128	3	0.30	3.00
6	4	256	5	3.54	3.44
7	3	256	2	5.41	3.75
8	3	256	3	3.57	3.92
9	5	128	5	0.54	3.93
10	3	256	5	5.83	5.29

**Table 8.3:** Performance of various power flow models using dataset configuration [C3].

Rank	$\mathcal{L}_{\mathcal{D}\mathcal{E}\mathcal{R}}$	$d$	$K$	Validation Loss (MAPE)	Test Loss (MAPE)
1	3	128	5	0.05	0.24 %
2	2	256	2	0.02	1.21 %
3	3	256	2	0.04	1.22 %
4	3	128	2	0.07	1.73 %
5	2	128	2	0.06	1.75 %
6	3	512	4	0.02	2.82 %
7	3	64	6	0.05	3.13 %
8	4	128	3	0.03	3.15 %
9	5	256	4	0.02	3.17 %
10	2	64	6	0.05	3.30 %

**Table 8.4:** Performance of various pre-trained models using dataset configuration [C1].

Rank	$\mathcal{L}_{DER}$	$d$	$K$	Validation Loss (MAPE)	Test Loss (MAPE)
1	3	128	5	0.01	0.01 %
2	2	256	2	0.02	0.22 %
3	3	256	2	0.03	0.23 %
4	3	128	2	0.05	0.52 %
5	2	128	2	0.05	0.53 %
6	3	512	4	0.04	0.66 %
7	4	128	3	0.04	0.75 %
8	2	64	6	0.02	0.82 %
9	3	64	6	0.02	0.83 %
10	3	128	3	0.03	0.83 %

**Table 8.5:** Performance of various pre-trained models using dataset configuration C2.

Rank	$\mathcal{L}_{DER}$	$d$	$K$	Validation Loss (MAPE)	Test Loss (MAPE)
1	2	256	2	0.01	0.23 %
2	3	256	2	0.03	0.34 %
3	4	128	3	0.03	0.49 %
4	3	128	2	0.04	0.77 %
5	2	128	2	0.16	0.93 %
6	4	256	4	0.03	0.96 %
7	4	512	4	0.01	1.00 %
8	3	64	6	0.01	1.03 %
9	2	64	6	0.01	1.04 %
10	3	512	4	0.03	1.13 %

**Table 8.6:** Performance of various pre-trained models using dataset configuration C3.

Rank	$\mathcal{L}_{DER}$	$d$	$K$	Validation Loss (MAPE)	Test Loss (MAPE)
1	3	128	3	0.05	1.57 %
2	3	512	3	0.03	2.40 %
3	4	512	3	0.02	2.62 %
4	3	256	4	0.03	2.74 %
5	2	256	4	0.03	2.76 %
6	5	128	3	0.04	2.86 %
7	5	512	3	0.02	2.95 %
8	4	256	3	0.03	3.17 %
9	3	256	3	0.02	3.87 %
10	2	256	5	0.02	4.04 %

**Table 8.7:** Performance of various cold-start models using dataset configuration  $[C1]$ .

Rank	$\mathcal{L}_{DER}$	$d$	$K$	Validation Loss (MAPE)	Test Loss (MAPE)
1	3	128	3	0.02	0.20 %
2	3	512	3	0.02	0.42 %
3	4	512	3	0.02	0.54 %
4	5	512	3	0.03	0.55 %
5	3	256	4	0.03	0.55 %
6	2	256	4	0.03	0.57 %
7	4	256	3	0.04	0.70 %
8	5	128	3	0.03	0.76 %
9	3	256	3	0.04	0.99 %
10	3	256	5	0.04	1.05 %

**Table 8.8:** Performance of various cold-start models using dataset configuration  $[C2]$ .

Rank	$\mathcal{L}_{\mathcal{DER}}$	$d$	$K$	Validation Loss (MAPE)	Test Loss (MAPE)
1	3	128	3	0.03	0.53 %
2	3	512	3	0.00	0.80 %
3	5	128	3	0.17	0.87 %
4	4	512	3	0.00	0.89 %
5	2	256	4	0.05	0.91 %
6	3	256	4	0.01	0.91 %
7	4	256	3	0.01	0.94 %
8	2	128	5	0.90	1.00 %
9	3	128	5	0.76	1.00 %
10	5	512	3	0.00	1.12 %

**Table 8.9:** Performance of various cold-start models using dataset configuration  $\boxed{\text{C3}}$ .



# Bibliography

---

[AAM22] Alireza Akrami, Salman Asif, and Hamed Mohsenian-Rad. **Sparse Tracking State Estimation for Low-Observable Power Distribution Systems Using D-PMUs**. *IEEE Transactions on Power Systems* 37:1 (Jan. 2022), 551–564. ISSN: 0885-8950, 1558-0679. DOI: 10.1109/TPWRS.2021.3094534. URL: <https://ieeexplore.ieee.org/document/9477195/> (visited on 09/09/2022) (see page 31).

[Ali+14] A. Alimardani, S. Zadkhast, J. Jatskevich, and E. Vaahedi. **Using smart meters in state estimation of distribution networks**. en. In: *2014 IEEE PES General Meeting / Conference & Exposition*. National Harbor, MD, USA: IEEE, July 2014, 1–5. ISBN: 978-1-4799-6415-4. DOI: 10.1109/PESGM.2014.6939801. URL: <http://ieeexplore.ieee.org/document/6939801/> (visited on 09/11/2022) (see page 31).

[And+11] Andreas Abart, Daniel Burnier, Benoit Bletterie, Matthias Stifter, Helfried Brunner, Andreas Lugmaier, and Alexander Schenk. **Power Snapshot Analysis: A new method for analyzing low voltage grids using a smart metering system**. en. In: vol. 21st. Frankfurt, June 2011 (see pages 31, 32).

[And+19] Andrew Banks, Ed Briggs, Ken Borgendale, and Rahul Gupta. *MQTT Version 5.0*. July 2019. URL: <https://docs.oasis-open.org/mqtt/mqtt/v5.0/mqtt-v5.0.html> (see page 32).

[Ang+16] Andrea Angioni, Tim Schlosser, Ferdinanda Ponci, and Antonello Monti. **Impact of Pseudo-Measurements From New Power Profiles on State Estimation in Low-Voltage Grids**. *IEEE Transactions on Instrumentation and Measurement* 65:1 (Jan. 2016), 70–77. ISSN: 0018-9456, 1557-9662. DOI: 10.1109/TIM.2015.2454673. URL: <http://ieeexplore.ieee.org/document/7312455/> (visited on 09/16/2024) (see pages 29, 33).

[Ara+19] M. A. Araujo, R. A. Flauzino, L. A. Moraes, F. A. S. Borges, and D. H. Spatti. **Decision Trees Applied to Fault Locations in Distribution Systems with Smart Meters**. en. In: *2019 IEEE International Conference on Environment and Electrical Engineering and 2019 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe)*. Genova, Italy: IEEE, June 2019, 1–6. ISBN: 978-1-72810-653-3. DOI: 10.1109/EEEIC.2019.8783511. URL: <https://ieeexplore.ieee.org/document/8783511/> (visited on 09/09/2022) (see page 31).

[ASB11] Ahmed E. B. Abu-Elanien, M. M. A. Salama, and Ray Bartnikas. **A Techno-Economic Method for Replacing Transformers**. en. *IEEE Transactions on Power Delivery* 26:2 (Apr. 2011), 817–829. ISSN: 0885-8977, 1937-4208. DOI: 10.1109/TPWRD.2010.2091289. URL: <http://ieeexplore.ieee.org/document/5674120/> (visited on 07/31/2025) (see page 26).

[bad23] badenova AG & Co. KG. *LoRaWAN: Einfach erklärt mit Anwendungen und Beispielen*. Jan. 2023. URL: <https://www.badenova.de/blog/lorawan-einfach-erklaert/> (visited on 02/16/2023) (see page 32).

[BAL17] Soheil Derafshi Beigvand, Hamdi Abdi, and Massimo La Scala. **Economic dispatch of multiple energy carriers**. en. *Energy* 138 (Nov. 2017), 861–872. ISSN: 03605442. doi: 10.1016/j.energy.2017.07.108. url: <https://linkinghub.elsevier.com/retrieve/pii/S0360544217312860> (visited on 02/06/2024) (see page 54).

[Bar+95] P.R. Barnes, J.W. Van Dyke, B.W. McConnell, S.M. Cohn, and S.L. Purucker. **The feasibility of replacing or upgrading utility distribution transformers during routine maintenance**. en. Tech. rep. ORNL-6804/R1, 82474. Apr. 1995, ORNL-6804/R1, 82474. doi: 10.2172/82474. url: <http://www.osti.gov/servlets/purl/82474-U3az7u/webviewable/> (visited on 07/31/2025) (see page 26).

[Bf23] Bundesministerium und für Wirtschaft und Klimaschutz. *Richtlinie für die Bundesförderung für effiziente Gebäude – Einzelmaßnahmen (BEG EM)*. de. Dec. 2023. url: <https://www.bundesanzeiger.de/pub/publication/TevdpcR9NeEp7m7RhbJ/content/TevdpcR9NeEp7m7RhbJ/BAnz0AT029.12.20230B1.pdf?inline> (see page 2).

[BJ09] Christian Blume and Wilfried Jakob. **GLEAM - General Learning Evolutionary Algorithm and Method : ein Evolutionärer Algorithmus und seine Anwendungen**. KIT Scientific Publishing, 2009. doi: 10.5445/KSP/1000013553. url: <https://publikationen.bibliothek.kit.edu/1000013553> (visited on 02/13/2024) (see pages 54, 55, 66).

[BN01] Aharon Ben-Tal and Arkadi Nemirovski. **Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications**. en. Society for Industrial and Applied Mathematics, Jan. 2001. ISBN: 978-0-89871-491-3 978-0-89871-882-9. doi: 10.1137/1.9780898718829. url: <http://pubs.siam.org/doi/book/10.1137/1.9780898718829> (visited on 07/30/2025) (see page 10).

[Bol+24] Suresh Bolusani, Mathieu Besançon, Ksenia Bestuzheva, Antonia Chmiela, João Dionísio, Tim Donkiewicz, Jasper van Doornmalen, Leon Eifler, Mohammed Ghannam, Ambros Gleixner, Christoph Graczyk, Katrin Halbig, Ivo Hedtke, Alexander Hoen, Christopher Hojny, Rolf van der Hulst, Dominik Kamp, Thorsten Koch, Kevin Kofler, Jürgen Lentz, Julian Manns, Gioni Mexi, Erik Mühmer, Marc E. Pfetsch, Franziska Schlösser, Felipe Serrano, Yuji Shinano, Mark Turner, Stefan Vigerske, Dieter Weninger, and Lixing Xu. **The SCIP Optimization Suite 9.0**. Technical Report. Optimization Online, Feb. 2024. url: <https://optimization-online.org/2024/02/the-scip-optimization-suite-9-0/> (see page 11).

[Bom+22] Rishi Bommasani et al. *On the Opportunities and Risks of Foundation Models*. en. arXiv:2108.07258 [cs]. July 2022. doi: 10.48550/arXiv.2108.07258. url: <http://arxiv.org/abs/2108.07258> (visited on 12/20/2024) (see page 92).

[BRZ19] Valentin Bolz, Johannes Rueß, and Andreas Zell. **Power Flow Approximation Based on Graph Convolutional Networks**. In: *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*. Boca Raton, FL, USA: IEEE, Dec. 2019, 1679–1686. ISBN: 978-1-72814-550-1. doi: 10.1109/ICMLA.2019.00274. url: <https://ieeexplore.ieee.org/document/8999165/> (visited on 01/19/2025) (see pages 90, 91).

[Byn+21] Michael L. Bynum, Gabriel A. Hackebeil, William E. Hart, Carl D. Laird, Bethany L. Nicholson, John D. Siirola, Jean-Paul Watson, and David L. Woodruff. **Pyomo-optimization modeling in python**. Third. Vol. 67. Springer Science & Business Media, 2021 (see pages 11, 56, 60).

[Cal+23] Katherine Calvin et al. **IPCC, 2023: Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (eds.)]**. IPCC, Geneva, Switzerland. Tech. rep. Edition: First. Intergovernmental Panel on Climate Change (IPCC), July 2023. doi: 10.59327/IPCC/AR6-9789291691647. URL: <https://www.ipcc.ch/report/ar6/syr/> (visited on 08/08/2025) (see page 1).

[Cho+16] Sridhar Chouhan, Deepak Tiwari, Hakan Inan, Sarika Khushalani-Solanki, and Ali Feliachi. **DER optimization to determine optimum BESS charge/discharge schedule using Linear Programming**. en. In: *2016 IEEE Power and Energy Society General Meeting (PESGM)*. Boston, MA, USA: IEEE, July 2016, 1–5. ISBN: 978-1-5090-4168-8. doi: 10.1109/PESGM.2016.7741576. URL: <http://ieeexplore.ieee.org/document/7741576/> (visited on 02/15/2024) (see page 54).

[Cos+11] E J Coster, J M A Myrzik, B Kruimer, and W L Kling. **Integration Issues of Distributed Generation in Distribution Grids**. en. *Proceedings of the IEEE* 99:1 (Jan. 2011), 28–39. ISSN: 0018-9219, 1558-2256. doi: 10.1109/JPROC.2010.2052776. URL: <http://ieeexplore.ieee.org/document/5565399/> (visited on 07/31/2024) (see page 33).

[Dah+20] Shweta Dahale, Hazhar Sufi Karimi, Kexing Lai, and Balasubramaniam Natarajan. **Sparsity Based Approaches for Distribution Grid State Estimation - A Comparative Study**. en. *IEEE Access* 8 (2020), 198317–198327. ISSN: 2169-3536. doi: 10.1109/ACCESS.2020.3035378. URL: <https://ieeexplore.ieee.org/document/9247106/> (visited on 08/12/2022) (see page 31).

[De +22] Steven De Jongh, Frederik Gielnik, Felicitas Mueller, Loris Schmit, Michael Suriyah, and Thomas Leibfried. **Physics-informed geometric deep learning for inference tasks in power systems**. en. *Electric Power Systems Research* 211 (Oct. 2022), 108362. ISSN: 03787796. doi: 10.1016/j.epsr.2022.108362. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0378779622005223> (visited on 06/04/2024) (see pages 18, 89–91, 106).

[Dem+25] Gökhan Demirel, Simon Grafenhorst, Kevin Förderer, and Veit Hagenmeyer. “Impact and Integration of Mini Photovoltaic Systems on Electric Power Distribution Grids.” en. In: *Lecture Notes in Computer Science*. ISSN: 0302-9743, 1611-3349. Cham: Springer Nature Switzerland, 2025, 247–265. ISBN: 978-3-031-74740-3 978-3-031-74741-0. doi: 10.1007/978-3-031-74741-0\_16. URL: [https://link.springer.com/10.1007/978-3-031-74741-0\\_16](https://link.springer.com/10.1007/978-3-031-74741-0_16) (visited on 07/21/2025) (see page 18).

[Di +18] Katia Gregio Di Santo, Silvio Giuseppe Di Santo, Renato Machado Monaro, and Marco Antonio Saidel. **Active demand side management for households in smart grids using optimization and artificial intelligence**. en. *Measurement* 115 (Feb. 2018), 152–161. ISSN: 02632241. doi: 10.1016/j.measurement.2017.10.010. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0263224117306346> (visited on 08/10/2025) (see page 7).

[Dim+14] P. Dimitriou, T. Leber, N. Nagele, L. Temmel, and C. Tessarek. **Voltage and frequency measuring plug: As part of Smart Grids metering system**. en. In: *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*. Venice, Italy: IEEE, Nov. 2014, 350–355. ISBN: 978-1-4799-4934-2. doi: 10.1109/SmartGridComm.2014.7007671. URL: <http://ieeexplore.ieee.org/document/7007671/> (visited on 08/11/2022) (see page 32).

[Du+17] Jian Du, Shanghang Zhang, Guanhong Wu, Jose M. F. Moura, and Soummya Kar. **Topology Adaptive Graph Convolutional Networks**. Version Number: 5. 2017. doi: 10.48550/ARXIV.1710.10370. URL: <https://arxiv.org/abs/1710.10370> (visited on 07/02/2024) (see pages 21, 91, 94).

[EB11] N. Etherden and M. H. J. Bollen. **Increasing the hosting capacity of distribution networks by curtailment of renewable energy resources**. en. In: *2011 IEEE Trondheim PowerTech*. Trondheim: IEEE, June 2011, 1–7. ISBN: 978-1-4244-8418-8 978-1-4244-8419-5 978-1-4244-8417-1. doi: 10.1109/PTC.2011.6019292. URL: <http://ieeexplore.ieee.org/document/6019292/> (visited on 02/27/2025) (see page 31).

[Eid+16] Cherrelle Eid, Paul Codani, Yannick Perez, Javier Reneses, and Rudi Hakvoort. **Managing electric flexibility from Distributed Energy Resources: A review of incentives for market design**. en. *Renewable and Sustainable Energy Reviews* 64 (Oct. 2016), 237–247. ISSN: 13640321. doi: 10.1016/j.rser.2016.06.008. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1364032116302222> (visited on 02/08/2024) (see page 9).

[Ela+20] Abdelfattah A. Eladl, Magda I. El-Afifi, Mohammed A. Saeed, and Magdi M. El-Saadawi. **Optimal operation of energy hubs integrated with renewable energy sources and storage devices considering CO<sub>2</sub> emissions**. en. *International Journal of Electrical Power & Energy Systems* 117 (May 2020), 105719. ISSN: 01420615. doi: 10.1016/j.ijepes.2019.105719. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0142061519314917> (visited on 02/06/2024) (see page 54).

[Ela+23] Abdelfattah A. Eladl, Magda I. El-Afifi, Magdi M. El-Saadawi, and Bishoy E. Sedhom. **A review on energy hubs: Models, methods, classification, applications, and future trends**. en. *Alexandria Engineering Journal* 68 (Apr. 2023), 315–342. ISSN: 11100168. doi: 10.1016/j.aej.2023.01.021. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1110016823000364> (visited on 01/08/2024) (see page 9).

[FPL20] Samar Fatima, Verner Püvi, and Matti Lehtonen. **Review on the PV Hosting Capacity in Distribution Networks**. en. *Energies* 13:18 (Sept. 2020), 4756. ISSN: 1996-1073. doi: 10.3390/en13184756. URL: <https://www.mdpi.com/1996-1073/13/18/4756> (visited on 07/19/2022) (see pages 26, 31).

[Gan+12] Tanuja Ganu, Deva P. Seetharam, Vijay Arya, Rajesh Kunnath, Jagabondhu Hazra, Saiful A Husain, Liyanage Chandratilake De Silva, and Shivkumar Kalyanaraman. **nPlug: a smart plug for alleviating peak loads**. en. In: *Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet*. Madrid Spain: ACM, May 2012, 1–10. ISBN: 978-1-4503-1055-0. doi: 10.1145/2208828.2208858. URL: <https://dl.acm.org/doi/10.1145/2208828.2208858> (visited on 01/20/2023) (see page 32).

[GBC16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. **Deep Learning**. <http://www.deeplearningbook.org>. MIT Press, 2016 (see pages 22, 98).

[Gei+07] Martin Geidl, Gaudenz Koeppel, Patrick Favre-Perrod, Bernd Klockl, Goran Andersson, and Klaus Frohlich. **Energy hubs for the future**. *IEEE Power and Energy Magazine* 5:1 (Jan. 2007), 24–30. ISSN: 1540-7977. DOI: [10.1109/MPAE.2007.264850](https://doi.org/10.1109/MPAE.2007.264850). URL: <http://ieeexplore.ieee.org/document/4042137/> (visited on 11/13/2023) (see pages 3, 7, 8, 54).

[Geo24] Philip George. “Finite Element Formulation for Heat Transfer.” en. In: *Introduction to Finite Element Analysis*. Ed. by S. Unnikrishnan Nair and S. Somanath. Singapore: Springer Nature Singapore, 2024, 191–230. ISBN: 978-981-19-7989-7. DOI: [10.1007/978-981-19-7989-7\\_7](https://doi.org/10.1007/978-981-19-7989-7_7). URL: [https://link.springer.com/10.1007/978-981-19-7989-7\\_7](https://link.springer.com/10.1007/978-981-19-7989-7_7) (visited on 08/13/2025) (see page 29).

[Ger25] German Federal Ministry for Economic Affairs and Climate Action. **Grids and infrastructure**. en (2025). URL: <https://www.bmwk.de/Redaktion/EN/Artikel/Energy/electricity-grids-of-the-future-01.html> (visited on 02/14/2025) (see page 12).

[GFH23] Simon Grafenhorst, Kevin Förderer, and Veit Hagenmeyer. *Poster Abstract: A three-phase electricity grid model of a single family house*. 12th DACH+ Conference on Energy Informatics (2023), Vienna, Austria, 4.–6. October 2023. 2023 (see pages 47, 113).

[GFH24a] Simon Grafenhorst, Kevin Förderer, and Veit Hagenmeyer. **Distribution grid monitoring based on feature propagation using smart plugs**. en. *Energy Informatics* 7:1 (Nov. 2024), 116. ISSN: 2520-8942. DOI: [10.1186/s42162-024-00427-y](https://doi.org/10.1186/s42162-024-00427-y). URL: <https://energyinformatics.springeropen.com/articles/10.1186/s42162-024-00427-y> (visited on 02/11/2025) (see pages 25, 34, 39, 45, 46, 48, 51, 56, 113).

[GFH24b] Simon Grafenhorst, Kevin Förderer, and Veit Hagenmeyer. “Distribution Grid Monitoring Based on Widely Available Smart Plugs.” en. In: *Energy Informatics*. Ed. by Bo Nørregaard Jørgensen, Luiz Carlos Pereira Da Silva, and Zheng Ma. Vol. 14468. Series Title: Lecture Notes in Computer Science. Cham: Springer Nature Switzerland, 2024, 195–212. ISBN: 978-3-031-48651-7 978-3-031-48652-4. DOI: [10.1007/978-3-031-48652-4\\_13](https://doi.org/10.1007/978-3-031-48652-4_13). URL: [https://link.springer.com/10.1007/978-3-031-48652-4\\_13](https://link.springer.com/10.1007/978-3-031-48652-4_13) (visited on 08/09/2024) (see pages 25, 27, 28, 37, 43, 48, 49, 113).

[GGK18] Etta Grover-Silva, Robin Girard, and George Kariniotakis. **Optimal sizing and placement of distribution grid connected battery systems through an SOCP optimal power flow algorithm**. en. *Applied Energy* 219 (June 2018), 385–393. ISSN: 03062619. DOI: [10.1016/j.apenergy.2017.09.008](https://doi.org/10.1016/j.apenergy.2017.09.008). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0306261917312813> (visited on 01/06/2024) (see pages 69, 71).

[Gha+22] Hosein Ghasemi, Jamshid Aghaei, Gevork B. Gharehpetian, Miadreza Shafie-Khah, and Joao P. S. Catalao. **Bi-Level Decomposition Approach for Coordinated Planning of an Energy Hub With Gas-Electricity Integrated Systems**. en. *IEEE Systems Journal* 16:1 (Mar. 2022), 1529–1539. ISSN: 1932-8184, 1937-9234, 2373-7816. DOI: [10.1109/JSYST.2021.3126592](https://doi.org/10.1109/JSYST.2021.3126592). URL: <https://ieeexplore.ieee.org/document/9632821/> (visited on 01/03/2024) (see page 71).

[Glo+22] John Duncan Glover, Thomas J. Overbye, Adam B. Birchfield, and Mulukutla S. Sarma. **Power System Analysis & Design: Si edition.** 7e. Boston: Cengage, 2022. ISBN: 978-0-357-67619-6 (see page 14).

[Gra+24a] Simon Grafenhorst, Gökhan Demirel, Kevin Förderer, and Veit Hagenmeyer. **Grid Aware Portfolio Optimization of a Multi-Energy DER.** en. In: *2024 IEEE Sustainable Power and Energy Conference (iSPEC)*. Kuching, Sarawak, Malaysia: IEEE, Nov. 2024, 245–250. ISBN: 9798350395075. doi: [10.1109/iSPEC59716.2024.10892444](https://doi.org/10.1109/iSPEC59716.2024.10892444). URL: <https://ieeexplore.ieee.org/document/10892444/> (visited on 03/07/2025) (see pages 69, 80, 84, 114).

[Gra+24b] Simon Grafenhorst, Jakob Ohm, Rafael Poppenborg, Kevin Förderer, and Veit Hagenmeyer. **Heuristic vs. Analytical Energy Hub Optimization: Design, Implementation, and Trade-Offs.** In: *2024 IEEE 12th International Conference on Smart Energy Grid Engineering (SEGE)*. Oshawa, ON, Canada: IEEE, Aug. 2024, 145–150. ISBN: 9798350377378. doi: [10.1109/SEGE62220.2024.10739589](https://doi.org/10.1109/SEGE62220.2024.10739589). URL: <https://ieeexplore.ieee.org/document/10739589/> (visited on 02/11/2025) (see pages 51, 58, 60–62, 64, 114).

[Gra+25] Simon Grafenhorst, Gökhan Demirel, Kevin Förderer, and Veit Hagenmeyer. **Augmented Pre-trained Graph Neural Networks for Grid-Supportive Flexibility Control.** en. In: *Proceedings of the 16th ACM International Conference on Future and Sustainable Energy Systems*. Rotterdam Netherlands: ACM, June 2025, 644–649. ISBN: 9798400711251. doi: [10.1145/3679240.3734663](https://doi.org/10.1145/3679240.3734663). URL: <https://dl.acm.org/doi/10.1145/3679240.3734663> (visited on 06/27/2025) (see pages 89, 106, 114).

[Gur23] Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual*. 2023. URL: <https://www.gurobi.com> (see pages 11, 65).

[HAI23] Abdul Hafeez, Rashid Alammari, and Atif Iqbal. **Utilization of EV Charging Station in Demand Side Management Using Deep Learning Method.** *IEEE Access* 11 (2023), 8747–8760. ISSN: 2169-3536. doi: [10.1109/ACCESS.2023.3238667](https://doi.org/10.1109/ACCESS.2023.3238667). URL: <https://ieeexplore.ieee.org/document/10024950/> (visited on 06/25/2025) (see page 7).

[Ham+24] Hendrik F. Hamann, Blazhe Gjorgiev, Thomas Brunschwiler, Leonardo S.A. Martins, Alban Puech, Anna Varbella, Jonas Weiss, Juan Bernabe-Moreno, Alexandre Blondin Massé, Seong Lok Choi, Ian Foster, Bri-Mathias Hodge, Rishabh Jain, Kibaek Kim, Vincent Mai, François Mirallès, Martin De Montigny, Octavio Ramos-Leaños, Hussein Suprême, Le Xie, El-Nasser S. Youssef, Arnaud Zinflou, Alexander Belyi, Ricardo J. Bessa, Bishnu Prasad Bhatarai, Johannes Schmude, and Stanislav Sobolevsky. **Foundation models for the electric power grid.** en. *Joule* 8:12 (Dec. 2024), 3245–3258. ISSN: 25424351. doi: [10.1016/j.joule.2024.11.002](https://doi.org/10.1016/j.joule.2024.11.002). URL: <https://linkinghub.elsevier.com/retrieve/pii/S2542435124004707> (visited on 08/12/2025) (see pages 2, 92).

[Hat+07] Nikos Hatziargyriou, Hiroshi Asano, Reza Iravani, and Chris Marnay. **Microgrids.** *IEEE Power and Energy Magazine* 5:4 (July 2007), 78–94. ISSN: 1540-7977. doi: [10.1109/MPAE.2007.376583](https://doi.org/10.1109/MPAE.2007.376583). URL: [http://ieeexplore.ieee.org/document/4263070/](https://ieeexplore.ieee.org/document/4263070/) (visited on 08/10/2025) (see page 7).

[Hus+18] Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. **Imitation Learning: A Survey of Learning Methods**. en. *ACM Computing Surveys* 50:2 (Mar. 2018), 1–35. ISSN: 0360-0300, 1557-7341. DOI: 10.1145/3054912. URL: <https://dl.acm.org/doi/10.1145/3054912> (visited on 12/17/2024) (see page 92).

[HWW11] William E Hart, Jean-Paul Watson, and David L Woodruff. **Pyomo: modeling and solving mathematical programs in Python**. *Mathematical Programming Computation* 3:3 (2011), 219–260 (see pages 11, 60).

[Int+24] International Renewable Energy Agency, COP28 Presidency, COP29 Presidency, Global Renewables Alliance, Ministry of Energy of the Republic of Azerbaijan, and Federal Government of Brazil. **Delivering on the UAE Consensus: Tracking progress toward tripling renewable energy capacity and doubling energy efficiency by 2030**. en (2024). ISSN: 978-92-9260-630-5. URL: [https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2024/Oct/IRENA\\_UAE\\_Consensus\\_2030\\_tripling\\_renewables\\_doubling\\_efficiency\\_2024.pdf](https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2024/Oct/IRENA_UAE_Consensus_2030_tripling_renewables_doubling_efficiency_2024.pdf) (visited on 01/14/2025) (see pages 1, 2).

[Jak+17] W. Jakob, J. Á. González Ordiano, N. Ludwig, R. Mikut, and V. Hagenmeyer. **Towards coding strategies for forecasting-based scheduling in smart grids and the energy lab 2.0**. en. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. Berlin Germany: ACM, July 2017, 1271–1278. ISBN: 978-1-4503-4939-0. DOI: 10.1145/3067695.3082481. URL: <https://dl.acm.org/doi/10.1145/3067695.3082481> (visited on 08/07/2025) (see pages 55, 64).

[Jan19] Janitza electronics GmbH. *The New Generation - the UMG PRO Series*. en. July 2019. URL: <https://www.janitza.de/files/download/leaflets/pro-serie/janitza-umg-pro-serie-en.pdf> (see page 35).

[Jay16] Dilan Jayaweera, ed. **Smart Power Systems and Renewable Energy System Integration**. en. Vol. 57. Studies in Systems, Decision and Control. Cham: Springer International Publishing, 2016. ISBN: 978-3-319-30427-4. DOI: 10.1007/978-3-319-30427-4. URL: <http://link.springer.com/10.1007/978-3-319-30427-4> (visited on 08/01/2025) (see page 16).

[JCL13] Ziyu Jia, Jiaxiong Chen, and Yuan Liao. **State estimation in distribution system considering effects of AMI data**. en. In: *2013 Proceedings of IEEE Southeastcon*. Jacksonville, FL, USA: IEEE, Apr. 2013, 1–6. ISBN: 978-1-4799-0053-4 978-1-4799-0052-7 978-1-4799-0051-0. DOI: 10.1109/SECON.2013.6567406. URL: <http://ieeexplore.ieee.org/document/6567406/> (visited on 09/06/2022) (see pages 31, 47).

[Jun22] Alexander Jung. **Machine Learning: The Basics**. en. Machine Learning: Foundations, Methodologies, and Applications. Singapore: Springer Nature Singapore, 2022. ISBN: 9789811681929 9789811681936. DOI: 10.1007/978-981-16-8193-6. URL: <https://link.springer.com/10.1007/978-981-16-8193-6> (visited on 07/31/2025) (see pages 22, 93).

[Kal00] Soteris A. Kalogirou. **Applications of artificial neural-networks for energy systems**. en. *Applied Energy* 67:1-2 (Sept. 2000), 17–35. ISSN: 03062619. DOI: 10.1016/S0306-2619(00)00005-2. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0306261900000052> (visited on 01/05/2025) (see page 91).

[KB17] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. en. arXiv:1412.6980 [cs]. Jan. 2017. URL: <http://arxiv.org/abs/1412.6980> (visited on 11/18/2024) (see pages 23, 94).

[KH22] Maman Ahmad Khan and Barry Hayes. **Smart Meter Based Two-Layer Distribution System State Estimation in Unbalanced MV/LV Networks**. en. *IEEE Transactions on Industrial Informatics* 18:1 (Jan. 2022), 688–697. ISSN: 1551-3203, 1941-0050. doi: [10.1109/TII.2021.3079267](https://doi.org/10.1109/TII.2021.3079267). URL: <https://ieeexplore.ieee.org/document/9432710/> (visited on 09/11/2022) (see page 31).

[Kos+24] Christoph Kost, Paul Müller, Schweiger, Jael Sepúlveda, Fluri, Verena, and Thomsen, Jessica. **Stromgestehungskosten Erneuerbare Energien**. de. *Fraunhofer-Institut für Solare Energiesysteme ISE* (July 2024). URL: [https://www.ise.fraunhofer.de/content/dam/ise/de/documents/publications/studies/DE2024\\_ISE\\_Studie\\_Stromgestehungskosten\\_Erneuerbare\\_Energien.pdf](https://www.ise.fraunhofer.de/content/dam/ise/de/documents/publications/studies/DE2024_ISE_Studie_Stromgestehungskosten_Erneuerbare_Energien.pdf) (see page 2).

[KP19] Nicolas Keriven and Gabriel Peyré. **Universal Invariant and Equivariant Graph Neural Networks**. In: *Advances in Neural Information Processing Systems*. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d' Alché-Buc, E. Fox, and R. Garnett. Vol. 32. Curran Associates, Inc., 2019. URL: [https://proceedings.neurips.cc/paper\\_files/paper/2019/file/ea9268cb43f55d1d12380fb6ea5bf572-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/ea9268cb43f55d1d12380fb6ea5bf572-Paper.pdf) (see page 33).

[LHL14] Yuan-Liang Lo, Shih-Che Huang, and Chan-Nan Lu. **Transformational Benefits of AMI Data in Transformer Load Modeling and Management**. en. *IEEE Transactions on Power Delivery* 29:2 (Apr. 2014), 742–750. ISSN: 0885-8977, 1937-4208. doi: [10.1109/TPWRD.2013.2280574](https://doi.org/10.1109/TPWRD.2013.2280574). URL: [http://ieeexplore.ieee.org/document/6607246/](https://ieeexplore.ieee.org/document/6607246/) (visited on 09/07/2022) (see pages 31, 32).

[Lia+22] Wenlong Liao, Birgitte Bak-Jensen, Jayakrishnan Radhakrishna Pillai, Yuelong Wang, and Yusen Wang. **A Review of Graph Neural Networks and Their Applications in Power Systems**. en. *Journal of Modern Power Systems and Clean Energy* 10:2 (2022), 345–360. ISSN: 2196-5625. doi: [10.35833/MPCE.2021.000058](https://doi.org/10.35833/MPCE.2021.000058). URL: <https://ieeexplore.ieee.org/document/9520300/> (visited on 08/05/2024) (see pages 89, 91).

[Lin+24] Nan Lin, Stavros Orfanoudakis, Nathan Ordonez Cardenas, Juan S. Giraldo, and Pedro P. Vergara. **PowerFlowNet: Power flow approximation using message passing Graph Neural Networks**. en. *International Journal of Electrical Power & Energy Systems* 160 (Sept. 2024), 110112. ISSN: 01420615. doi: [10.1016/j.ijepes.2024.110112](https://doi.org/10.1016/j.ijepes.2024.110112). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0142061524003338> (visited on 11/08/2024) (see pages 18, 20, 89–91, 94, 95, 106, 108).

[Liu+12] Xiaohu Liu, Andreas Aichhorn, Liming Liu, and Hui Li. **Coordinated Control of Distributed Energy Storage System With Tap Changer Transformers for Voltage Rise Mitigation Under High Photovoltaic Penetration**. en. *IEEE Transactions on Smart Grid* 3:2 (June 2012), 897–906. ISSN: 1949-3053, 1949-3061. doi: [10.1109/TSG.2011.2177501](https://doi.org/10.1109/TSG.2011.2177501). URL: [http://ieeexplore.ieee.org/document/6145728/](https://ieeexplore.ieee.org/document/6145728/) (visited on 02/08/2024) (see page 9).

[Liu+24] Jiawei Liu, Cheng Yang, Zhiyuan Lu, Junze Chen, Yibo Li, Mengmei Zhang, Ting Bai, Yuan Fang, Lichao Sun, Philip S. Yu, and Chuan Shi. *Towards Graph Foundation Models: A Survey and Beyond*. \_eprint: 2310.11829. 2024. url: <https://arxiv.org/abs/2310.11829> (see page 90).

[LLW17] Rasmus Luthander, David Lingfors, and Joakim Widén. **Large-scale integration of photovoltaic power in a distribution grid using power curtailment and energy storage**. en. *Solar Energy* 155 (Oct. 2017), 1319–1325. issn: 0038092X. doi: 10.1016/j.solener.2017.07.083. url: <https://linkinghub.elsevier.com/retrieve/pii/S0038092X17306680> (visited on 01/10/2023) (see page 12).

[LoR] LoRa Alliance. *What is LoRaWAN®?* url: <https://lora-alliance.org/about-lorawan/> (visited on 08/11/2025) (see page 32).

[LRM17] Bei Li, Robin Roche, and Abdellatif Miraoui. **Microgrid sizing with combined evolutionary algorithm and MILP unit commitment**. en. *Applied Energy* 188 (Feb. 2017), 547–562. issn: 03062619. doi: 10.1016/j.apenergy.2016.12.038. url: <https://linkinghub.elsevier.com/retrieve/pii/S0306261916318013> (visited on 02/13/2024) (see page 54).

[LVZ14] M. La Scala, A. Vaccaro, and A.F. Zobaa. **A goal programming methodology for multiobjective optimization of distributed energy hubs operation**. en. *Applied Thermal Engineering* 71:2 (Oct. 2014), 658–666. issn: 13594311. doi: 10.1016/j.applthermaleng.2013.10.031. url: <https://linkinghub.elsevier.com/retrieve/pii/S1359431113007400> (visited on 01/08/2024) (see page 71).

[M A74] M. A. Stephens. **EDF Statistics for Goodness of Fit and Some Comparisons**. en. *Journal of the American Statistical Association* (1974), 9 (see page 43).

[Man+12] Efthymios Manitsas, Ravindra Singh, Bikash C. Pal, and Goran Strbac. **Distribution System State Estimation Using an Artificial Neural Network Approach for Pseudo Measurement Modeling**. *IEEE Transactions on Power Systems* 27:4 (Nov. 2012), 1888–1896. issn: 0885-8950, 1558-0679. doi: 10.1109/TPWRS.2012.2187804. url: <http://ieeexplore.ieee.org/document/6176289/> (visited on 09/16/2024) (see pages 29, 33, 34).

[Man+14] Kebina Manandhar, Xiaojun Cao, Fei Hu, and Yao Liu. **Detection of Faults and Attacks Including False Data Injection Attack in Smart Grid Using Kalman Filter**. en. *IEEE Transactions on Control of Network Systems* 1:4 (Dec. 2014), 370–379. issn: 2325-5870. doi: 10.1109/TCNS.2014.2357531. url: <http://ieeexplore.ieee.org/document/6897944/> (visited on 10/07/2024) (see page 33).

[Mar22] Martin Bäckman. *Smart Homes and Home Automation, Smart Buildings 9th Edition*. en. Apr. 2022 (see pages 2, 26).

[Mar24] Martin Bäckman. *Smart Homes and Home Automation*. en. Oct. 2024. url: <https://media.berginsight.com/2024/10/23231726/bi-sh11-ps.pdf> (see page 26).

[Mas+17] Salman Mashayekh, Michael Stadler, Gonçalo Cardoso, and Miguel Heleno. **A mixed integer linear programming approach for optimal DER portfolio, sizing, and placement in multi-energy microgrids.** en. *Applied Energy* 187 (Feb. 2017), 154–168. ISSN: 03062619. DOI: 10.1016/j.apenergy.2016.11.020. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0306261916316051> (visited on 07/18/2023) (see page 71).

[Mat+12] E. Matallanas, M. Castillo-Cagigal, A. Gutiérrez, F. Monasterio-Huelin, E. Caamaño-Martín, D. Masa, and J. Jiménez-Leube. **Neural network controller for Active Demand-Side Management with PV energy in the residential sector.** en. *Applied Energy* 91:1 (Mar. 2012), 90–97. ISSN: 03062619. DOI: 10.1016/j.apenergy.2011.09.004. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0306261911005630> (visited on 08/10/2025) (see page 7).

[MBE20] Enock Mulenga, Math H.J. Bollen, and Nicholas Etherden. **A review of hosting capacity quantification methods for photovoltaics in low-voltage distribution grids.** en. *International Journal of Electrical Power & Energy Systems* 115 (Feb. 2020), 105445. ISSN: 01420615. DOI: 10.1016/j.ijepes.2019.105445. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0142061519306490> (visited on 09/12/2022) (see page 26).

[MBH23] Marija Marković, Matthew Bossart, and Bri-Mathias Hodge. **Machine learning for modern power distribution systems: Progress and perspectives.** en. *Journal of Renewable and Sustainable Energy* 15:3 (May 2023), 032301. ISSN: 1941-7012. DOI: 10.1063/5.0147592. URL: <https://pubs.aip.org/jrse/article/15/3/032301/2900695> Machine - learning - for - modern - power - distribution (visited on 06/04/2024) (see page 33).

[Mei+20] Steffen Meinecke, Džanan Sarajlić, Simon Ruben Drauz, Annika Klettke, Lars-Peter Lauven, Christian Rehtanz, Albert Moser, and Martin Braun. **SimBench—A Benchmark Dataset of Electric Power Systems to Compare Innovative Solutions Based on Power Flow Analysis.** en. *Energies* 13:12 (June 2020), 3290. ISSN: 1996-1073. DOI: 10.3390/en13123290. URL: <https://www.mdpi.com/1996-1073/13/12/3290> (visited on 01/18/2023) (see pages 44, 93).

[Mel+21] Gustavo Costa Gomes de Melo, Igor Cavalcante Torres, Ícaro Bezerra Queiroz de Araújo, Davi Bibiano Brito, and Erick de Andrade Barboza. **A Low-Cost IoT System for Real-Time Monitoring of Climatic Variables and Photovoltaic Generation for Smart Grid Application.** en. *Sensors* 21:9 (May 2021), 3293. ISSN: 1424-8220. DOI: 10.3390/s21093293. URL: <https://www.mdpi.com/1424-8220/21/9/3293> (visited on 09/09/2022) (see page 32).

[Mif+22] Bardhyl Miftari, Mathias Berger, Hatim Djelassi, and Damien Ernst. **GBOML: Graph-Based Optimization ModelingLanguage.** *Journal of Open Source Software* 7:72 (Apr. 2022), 4158. ISSN: 2475-9066. DOI: 10.21105/joss.04158. URL: <https://joss.theoj.org/papers/10.21105/joss.04158> (visited on 01/03/2024) (see page 11).

[Mit+15] Rajendu Mitra, Ramachandra Kota, Sambaran Bandyopadhyay, Vijay Arya, Brian Sullivan, Richard Mueller, Heather Storey, and Gerard Labut. **Voltage Correlations in Smart Meter Data.** en. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. Sydney NSW Australia: ACM, Aug. 2015, 1999–2008. ISBN: 978-1-4503-3664-2. DOI: 10.1145/2783258.2788594. URL:

<https://dl.acm.org/doi/10.1145/2783258.2788594> (visited on 02/26/2025) (see pages 2, 12).

[Naj+22] Arsalan Najafi, Mahdi Pourakbari-Kasmaei, Michal Jasinski, Matti Lehtonen, and Zbigniew Leonowicz. **A medium-term hybrid IGDT-Robust optimization model for optimal self scheduling of multi-carrier energy systems**. en. *Energy* 238 (Jan. 2022), 121661. ISSN: 03605442. doi: 10.1016/j.energy.2021.121661. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0360544221019095> (visited on 02/15/2024) (see page 54).

[Ner+25] Christian H. Nerowski, Simon Grafenhorst, Susanne Schmitt, Martin Lindner, Kevin Förderer, Christian Rehtanz, and Veit Hagenmeyer. **A modular Architecture for Congestion Management using Cross Voltage Level Power Flow Controllers and Energy Hubs**. In: *ETG-Fb. 176: ETG Kongress 2025 : Voller Energie – heute und morgen. 21. – 22. Mai 2025 in Kassel*. ETG Kongress. 2025 (Kassel, Deutschland, May 21–22, 2025). 37.12.03; LK 01. VDE Verlag, 2025. ISBN: 978-3-8007-6494-5 (see pages 8, 52).

[Nev+17] Neven Duić, Nedeljko Štefanić, Zoran Lulić, Goran Krajačić, Tomislav Pukšec, and Tomislav Novose. *EU28 fuel prices for 2015, 2030 and 2050*. en. Nov. 2017. URL: [https://heatroadmap.eu/wp-content/uploads/2020/01/HRE4\\_D6.1-Future-fuel-price-review.pdf](https://heatroadmap.eu/wp-content/uploads/2020/01/HRE4_D6.1-Future-fuel-price-review.pdf) (see page 2).

[NG25] National Aeronautics and Space Administration and Goddard Institute for Space Studies. *GISS Surface Temperature Analysis (v4), Annual Mean Temperature Change over Land and over Ocean*. en. Aug. 2025. URL: [https://data.giss.nasa.gov/gistemp/graphs\\_v4/](https://data.giss.nasa.gov/gistemp/graphs_v4/) (visited on 08/08/2025) (see page 1).

[PL17] Anggoro Primadianto and Chan-Nan Lu. **A Review on Distribution System State Estimation**. en. *IEEE TRANSACTIONS ON POWER SYSTEMS* 32:5 (2017), 9 (see page 31).

[Pop+21] Rafael Poppenborg, Johannes Ruf, Malte Chlost, Jianlei Liu, Christian Hotz, Clemens Düpmeier, Thomas Kolb, and Veit Hagenmeyer. **Energy hub gas: a multi-domain system modelling and co-simulation approach**. en. In: *Proceedings of the 9th Workshop on Modeling and Simulation of Cyber-Physical Energy Systems*. Virtual Event: ACM, May 2021, 1–7. ISBN: 978-1-4503-8608-1. doi: 10.1145/3470481.3472712. URL: <https://dl.acm.org/doi/10.1145/3470481.3472712> (visited on 12/19/2023) (see page 72).

[Pop+22] Rafael Poppenborg, Malte Chlost, Johannes Ruf, Christian Hotz, Clemens Dupmeier, Thomas Kolb, and Veit Hagenmeyer. **Energy Hub Gas: A Modular Setup for the Evaluation of Local Flexibility and Renewable Energy Carriers Provision**. en. In: *2022 IEEE 10th International Conference on Smart Energy Grid Engineering (SEGE)*. Oshawa, ON, Canada: IEEE, Aug. 2022, 33–41. ISBN: 978-1-66549-930-9. doi: 10.1109/SEGE55279.2022.9889751. URL: <https://ieeexplore.ieee.org/document/9889751/> (visited on 12/19/2023) (see page 72).

[Pop+23a] Rafael Poppenborg, Katharina Beisswanger, Christian Hotz, Kevin Förderer, Thomas Kolb, and Veit Hagenmeyer. **Dynamic Mapping for Evolutionary Algorithm Based Optimization of Energy Hub Gas Scheduling**. en. In: *2023 IEEE 11th International Conference on Smart Energy Grid Engineering (SEGE)*. Oshawa, ON, Canada: IEEE,

Aug. 2023, 206–211. ISBN: 9798350340716. DOI: [10.1109/SEGE59172.2023.10274571](https://doi.org/10.1109/SEGE59172.2023.10274571). URL: <https://ieeexplore.ieee.org/document/10274571/> (visited on 01/29/2024) (see pages 55, 57, 58, 64, 66).

[Pop+23b] Rafael Poppenborg, Kaleb Phipps, Hatem Khalloof, Kevin Förderer, Ralf Mikut, and Veit Hagenmeyer. **Dynamic Chromosome Interpretation in Evolutionary Algorithms for Distributed Energy Resources Scheduling**. en. In: *Proceedings of the Companion Conference on Genetic and Evolutionary Computation*. Lisbon Portugal: ACM, July 2023, 755–758. ISBN: 9798400701207. DOI: [10.1145/3583133.3590666](https://doi.org/10.1145/3583133.3590666). URL: <https://dl.acm.org/doi/10.1145/3583133.3590666> (visited on 08/07/2025) (see page 66).

[Pop24] Rafael Poppenborg. *Optimized Energy Hub Scheduling Using Evolutionary Algorithms*. en. Medium: PDF. 2024. DOI: [10.5445/IR/1000177487](https://doi.org/10.5445/IR/1000177487). URL: <https://publikationen.bibliothek.kit.edu/1000177487> (visited on 08/04/2025) (see page 54).

[PR02] V.Leonardo Paucar and Marcos J Rider. **Artificial neural networks for solving the power flow problem in electric power systems**. en. *Electric Power Systems Research* 62:2 (June 2002), 139–144. ISSN: 03787796. DOI: [10.1016/S0378-7796\(02\)00030-5](https://doi.org/10.1016/S0378-7796(02)00030-5). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0378779602000305> (visited on 01/06/2025) (see pages 89, 91).

[Qui+15] Franklin L. Quilumba, Wei-Jen Lee, Heng Huang, David Y. Wang, and Robert L. Szabados. **Using Smart Meter Data to Improve the Accuracy of Intraday Load Forecasting Considering Customer Behavior Similarities**. en. *IEEE Transactions on Smart Grid* 6:2 (Mar. 2015), 911–918. ISSN: 1949-3053, 1949-3061. DOI: [10.1109/TSG.2014.2364233](https://doi.org/10.1109/TSG.2014.2364233). URL: [http://ieeexplore.ieee.org/document/6945384/](https://ieeexplore.ieee.org/document/6945384/) (visited on 09/12/2022) (see page 31).

[Ros+21] Emanuele Rossi, Henry Kenlay, Maria I. Gorinova, Benjamin Paul Chamberlain, Xiaowen Dong, and Michael Bronstein. *On the Unreasonable Effectiveness of Feature propagation in Learning on Graphs with Missing Node Features*. Version Number: 3. 2021. DOI: [10.48550/ARXIV.2111.12128](https://doi.org/10.48550/ARXIV.2111.12128). URL: <https://arxiv.org/abs/2111.12128> (visited on 06/04/2024) (see pages 3, 29, 33, 34, 38, 40).

[Sam+11] Kamalanath Samarakoon, Jianzhong Wu, Janaka Ekanayake, and Nick Jenkins. **Use of delayed smart meter measurements for distribution state estimation**. en. In: *2011 IEEE Power and Energy Society General Meeting*. San Diego, CA: IEEE, July 2011, 1–6. ISBN: 978-1-4577-1000-1 978-1-4577-1002-5. DOI: [10.1109/PES.2011.6039384](https://doi.org/10.1109/PES.2011.6039384). URL: <https://ieeexplore.ieee.org/document/6039384/> (visited on 09/06/2022) (see pages 31, 47).

[San+17] Ruben Sanchez, Florin Iov, Mohammed Kemal, Maria Stefan, and Rasmus Olsen. **Observability of low voltage grids: Actual DSOs challenges and research questions**. en. In: *2017 52nd International Universities Power Engineering Conference (UPEC)*. Heraklion: IEEE, Aug. 2017, 1–6. ISBN: 978-1-5386-2344-2. DOI: [10.1109/UPEC.2017.8232008](https://doi.org/10.1109/UPEC.2017.8232008). URL: [http://ieeexplore.ieee.org/document/8232008/](https://ieeexplore.ieee.org/document/8232008/) (visited on 09/11/2022) (see pages 31, 47).

[Sas+20] Susanne Sass, Timm Faulwasser, Dinah Elena Hollermann, Chrysoula Dimitra Kappatou, Dominique Sauer, Thomas Schütz, David Yang Shu, André Bardow, Lutz Gröll, Veit Hagenmeyer, Dirk Müller, and Alexander Mitsos. **Model compendium, data, and optimization benchmarks for sector-coupled energy systems**. en. *Computers & Chemical Engineering* 135 (Apr. 2020), 106760. ISSN: 00981354. DOI: 10.1016/j.compchemeng.2020.106760. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0098135419310683> (visited on 02/20/2024) (see page 65).

[Sca+09] F. Scarselli, M. Gori, Ah Chung Tsoi, M. Hagenbuchner, and G. Monfardini. **The Graph Neural Network Model**. en. *IEEE Transactions on Neural Networks* 20:1 (Jan. 2009), 61–80. ISSN: 1045-9227, 1941-0093. DOI: 10.1109/TNN.2008.2005605. URL: <http://ieeexplore.ieee.org/document/4700287/> (visited on 01/05/2025) (see page 91).

[Sch15] Jürgen Schmidhuber. **Deep learning in neural networks: An overview**. en. *Neural Networks* 61 (Jan. 2015), 85–117. ISSN: 08936080. DOI: 10.1016/j.neunet.2014.09.003. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0893608014002135> (visited on 12/17/2024) (see page 92).

[Sha+17] Chengcheng Shao, Xifan Wang, Mohammad Shahidehpour, Xiuli Wang, and Biyang Wang. **An MILP-Based Optimal Power Flow in Multicarrier Energy Systems**. en. *IEEE Transactions on Sustainable Energy* 8:1 (Jan. 2017), 239–248. ISSN: 1949-3029, 1949-3037. DOI: 10.1109/TSTE.2016.2595486. URL: <http://ieeexplore.ieee.org/document/7523971/> (visited on 01/08/2024) (see page 71).

[Sha+21] Ruibo Shang, Kevin Hoffer-Hawlik, Fei Wang, Guohai Situ, and Geoffrey P. Luke. **Two-step training deep learning framework for computational imaging without physics priors**. en. *Optics Express* 29:10 (May 2021), 15239. ISSN: 1094-4087. DOI: 10.1364/OE.424165. URL: <https://opg.optica.org/abstract.cfm?URI=oe-29-10-15239> (visited on 12/17/2024) (see page 92).

[Son+17] Yonghua Song, Jin Lin, Ming Tang, and Shufeng Dong. **An Internet of Energy Things Based on Wireless LPWAN**. en. *Engineering* 3:4 (Aug. 2017), 460–466. ISSN: 20958099. DOI: 10.1016/J.ENG.2017.04.011. URL: <https://linkinghub.elsevier.com/retrieve/pii/S2095809917306057> (visited on 09/08/2022) (see page 36).

[SP89] D.J. Slobajic and Y.-H. Pao. **Artificial neural-net based dynamic security assessment for electric power systems**. *IEEE Transactions on Power Systems* 4:1 (Feb. 1989), 220–228. ISSN: 08858950. DOI: 10.1109/59.32481. URL: <http://ieeexplore.ieee.org/document/32481/> (visited on 01/05/2025) (see page 91).

[SPJ10] R. Singh, B.C. Pal, and R.A. Jabr. **Statistical Representation of Distribution System Loads Using Gaussian Mixture Model**. *IEEE Transactions on Power Systems* 25:1 (Feb. 2010), 29–37. ISSN: 0885-8950, 1558-0679. DOI: 10.1109/TPWRS.2009.2030271. URL: <http://ieeexplore.ieee.org/document/5298967/> (visited on 09/16/2024) (see pages 29, 33, 34).

[SS17] Bindeshwar Singh and Janmejay Sharma. **A review on distributed generation planning**. en. *Renewable and Sustainable Energy Reviews* 76 (Sept. 2017), 529–544. ISSN: 13640321. DOI: 10.1016/j.rser.2017.03.034. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1364032117303568> (visited on 05/05/2025) (see page 90).

[SSU20] Indro Spinelli, Simone Scardapane, and Aurelio Uncini. **Missing data imputation with adversarially-trained graph convolutional networks**. en. *Neural Networks* 129 (Sept. 2020), 249–260. ISSN: 08936080. DOI: 10.1016/j.neunet.2020.06.005. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0893608020302185> (visited on 07/30/2024) (see pages 34, 40, 41, 44, 45, 48, 49).

[SV23] Oded Schwartz and Noa Vaknin. **Pebbling Game and Alternative Basis for High Performance Matrix Multiplication**. en. *SIAM Journal on Scientific Computing* 45:6 (Dec. 2023), C277–C303. ISSN: 1064-8275, 1095-7197. DOI: 10.1137/22M1502719. URL: <https://pubs.siam.org/doi/10.1137/22M1502719> (visited on 09/24/2024) (see page 41).

[SW65] S S Shapiro and M B Wilk. **An Analysis of Variance Test for Normality (Complete Samples)**. en. *Biometrika* 52:3/4 (1965), 22 (see page 43).

[Swa+21] Swantje Gährs, Julika Weiß, Hannes Bluhm, Elisa Dunkelberg, Jannes Katner, and Institut für ökologische Wirtschaftsforschung, Berlin. **Erkenntnisse zu Umweltwirkungen von Smart Metern: Erfahrungen aus dem Einsatz von Smart Metern in Europa**. de. *Ressortforschungsplan des Bundesministerium für Umwelt, Naturschutz und nukleare Sicherheit* CLIMATE CHANGE 34/2021 (2021) (see page 12).

[Tas22] Tasmota. *Tasmota Components*. en. May 2022. URL: <https://tasmota.github.io/docs/Components/%5C#tasmota> (see pages 35, 42).

[TF17] Fernanda C. L. Trindade and Walmir Freitas. **Low Voltage Zones to Support Fault Location in Distribution Systems With Smart Meters**. en. *IEEE Transactions on Smart Grid* 8:6 (Nov. 2017), 2765–2774. ISSN: 1949-3053, 1949-3061. DOI: 10.1109/TSG.2016.2538268. URL: <http://ieeexplore.ieee.org/document/7440854/> (visited on 09/09/2022) (see page 31).

[Thu+18] Leon Thurner, Alexander Scheidler, Florian Schafer, Jan-Hendrik Menke, Julian Dollichon, Friederike Meier, Steffen Meinecke, and Martin Braun. **Pandapower—An Open-Source Python Tool for Convenient Modeling, Analysis, and Optimization of Electric Power Systems**. en. *IEEE Transactions on Power Systems* 33:6 (Nov. 2018), 6510–6521. ISSN: 0885-8950, 1558-0679. DOI: 10.1109/TPWRS.2018.2829021. URL: <https://ieeexplore.ieee.org/document/8344496/> (visited on 09/11/2022) (see pages 18, 73).

[Tzi+20] Efthymios Tzinis, Shrikant Venkataramani, Zhepei Wang, Cem Subakan, and Paris Smaragdis. **Two-Step Sound Source Separation: Training On Learned Latent Targets**. en. In: *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Barcelona, Spain: IEEE, May 2020, 31–35. ISBN: 978-1-5090-6631-5. DOI: 10.1109/ICASSP40776.2020.9054172. URL: <https://ieeexplore.ieee.org/document/9054172/> (visited on 12/17/2024) (see page 92).

[VMJ21] Nicolai Bo Vanting, Zheng Ma, and Bo Nørregaard Jørgensen. **A scoping review of deep neural networks for electric load forecasting**. en. *Energy Informatics* 4:S2 (Sept. 2021), 49. ISSN: 2520-8942. DOI: 10.1186/s42162-021-00148-6. URL: <https://energyinformatics.springeropen.com/articles/10.1186/s42162-021-00148-6> (visited on 10/07/2024) (see page 32).

[Wan+17] Jianxiao Wang, Haiwang Zhong, Ziming Ma, Qing Xia, and Chongqing Kang. **Review and prospect of integrated demand response in the multi-energy system**. en. *Applied Energy* 202 (Sept. 2017), 772–782. ISSN: 03062619. doi: 10.1016/j.apenergy.2017.05.150. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0306261917306992> (visited on 02/08/2024) (see page 9).

[Wan+19] Yi Wang, Qixin Chen, Tao Hong, and Chongqing Kang. **Review of Smart Meter Data Analytics: Applications, Methodologies, and Challenges**. en. *IEEE Transactions on Smart Grid* 10:3 (May 2019), 3125–3148. ISSN: 1949-3053, 1949-3061. doi: 10.1109/TSG.2018.2818167. URL: <https://ieeexplore.ieee.org/document/8322199/> (visited on 09/11/2022) (see page 31).

[Wid+10] Joakim Widén, Ewa Wäckelgård, Jukka Paatero, and Peter Lund. **Impacts of distributed photovoltaics on network voltages: Stochastic simulations of three Swedish low-voltage distribution grids**. en. *Electric Power Systems Research* 80:12 (Dec. 2010), 1562–1571. ISSN: 03787796. doi: 10.1016/j.epsr.2010.07.007. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0378779610001707> (visited on 03/30/2023) (see page 12).

[Wor24] World Bank. **State and Trends of Carbon Pricing 2024**. en. Washington, DC: World Bank, 2024. ISBN: 978-1-4648-2127-1. URL: <https://hdl.handle.net/10986/41544> (see page 2).

[Zha+23] Jinxian Zhang, Junbo Zhao, Jing Yang, and Junhui Zhao. **Deep Multi-Fidelity Bayesian Data Fusion for Probabilistic Distribution System Voltage Estimation With High Penetration of PVs**. en. *IEEE Transactions on Power Systems* (2023), 1–12. ISSN: 0885-8950, 1558-0679. doi: 10.1109/TPWRS.2023.3295795. URL: <https://ieeexplore.ieee.org/document/10184461/> (visited on 09/18/2024) (see pages 29, 33).

[Zho+22] Xiaoqing Zhong, Weifeng Zhong, Yi Liu, Chao Yang, and Shengli Xie. **Optimal energy management for multi-energy multi-microgrid networks considering carbon emission limitations**. en. *Energy* 246 (May 2022), 123428. ISSN: 03605442. doi: 10.1016/j.energy.2022.123428. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0360544222003310> (visited on 02/17/2024) (see pages 9, 54).