Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Federal Department of the Environment,
Transport, Energy and Communications DETEC

Swiss Federal Office of Energy
Energy Research and Cleantech

Final report from 16 December 2025

LASAGNE

Digital framework for smart grid and renewable energies

Publisher:
Swiss Federal Office of Energy SFOE
Energy Research and Cleantech
CH-3003 Berne
www.energy-research.ch

Subsidy recipients:
HES-SO – HEPIA, 4, rue de la Prairie, 1202 Geneva
HES-SO VS, 3, rue de Technopôle, 3960 Sierre
UNIGE – Centre Universitaire d'Informatique, Battelle Bât. A, 7, rte de Drize, 1227 Carouge
CLEMAP AG, Flurstrasse 30, 8048 Zürich

Authors:
Nabil Abdennadher, HES-SO – HEPIA, nabil.abdennadher@hesge.ch
Jérôme Chételat, HES-SO – HEPIA
Raoul Dupuis, HES-SO – HEPIA
Mohsen Salimi, HES-SO – HEPIA
John White, HES-SO – HEPIA
Emmanuel Fragnière, HES-SO//VS, emmanuel.fragniere@hevs.ch
Giovanna Di Marzo Serugendo, UNIGE ‑ CUI, giovanna.dimarzo@unige.ch
Pascal Kienast, CLEMAP AG, pascal@clemap.ch
Gino Agbomemewa, CLEMAP AG
Fabian Weiersmüller, CLEMAP AG
Carlo Fischione, KTH
Josefin Danielsson, ElectriCITY

SFOE project coordinator:
Michael Moser, michael.moser@bfe.admin.ch

SFOE contract number: SI/502357-01

digitaL frAmework for SmArt Grid and reNewable Energies (LASAGNE)

Transnational ERA-Net SES project (108767)

Final technical Report

December 2025

Chapter 1: Introduction

This report details the achievements of the transnational ERA-Net SES project (108767), entitled "digitaL frAmework for SmArt Grid and reNewable Energie" (LASAGNE). The consortium is composed of 6 partners: three from Sweden (RECAP/TVINN, KTH and ElectriCITY) and three from Switzerland (HES-SO, UNIGE and CLEMAP).

The energy transition calls for smart meters to support functions other than measuring and transmitting electricity consumption and production. Apart from monitoring home appliances' power cycles, they will need to gather data about other households so as to predict/plan local consumption/production. These novel smart meters (Grid Edge Devices, "GED"), need to be managed in a collaborative fashion. In this context, the goal of LASAGNE is twofold:

1. Develop a Grid Edge Device (GED) enhanced with collaborative AI algorithms which are the foundation to build context-aware energy applications. The GED to develop relies on the devices produced and marketed by CLEMAP.
2. Develop and deploy a digital framework for smart grid and renewable energies and involve four stakeholders: (1) a system integrator (Distribution Systems Operators: DSO, Local Energy Communities: LECs, municipalities, buildings owners, etc.), (2) independent software vendors, (3) edge equipment vendors and (4) need-owners.

The framework is deployed in two pilots in Switzerland and one Living Lab in Sweden. The project also offers a "vision/concept" for developing a marketplace to empower this framework by features allowing stakeholders to implement their financial/technical interactions.

LASAGNE targets flexibility applications (Fig. 1.1). Flexibility is defined here as the ability to shift/shed the peak load from the demand and/or supply side. This definition applies to different levels: grids, microgrids (Local Energy communities: LECs) and buildings.

The flexibility energy applications are assumed to achieve several goals such as minimizing energy costs for Electric Vehicle (EV) charging, guaranteeing the timely and complete charging of vehicles, harnessing e-vehicles and charging stations as dynamic load suppliers to fortify the grid, safeguarding infrastructure performance, optimising localised energy exchanges, ensuring voltage stability, and enabling the efficient aggregation of distributed energy resources. Section 6 of the last progress report (2024)[1] described some of the applications targeted by LASAGNE.

Flexibility applications rely on forecasting tools that predict the consumption/production of energy. The prediction period (horizon) depends on the application itself. For security and privacy reasons, these forecast modules are deployed close to the data (edge-based digital infrastructure in Fig. 1.1).

---

[1] https://lsds.hesge.ch/wp-content/uploads/2025/10/LASAGNE_ProgressReport2.pdf

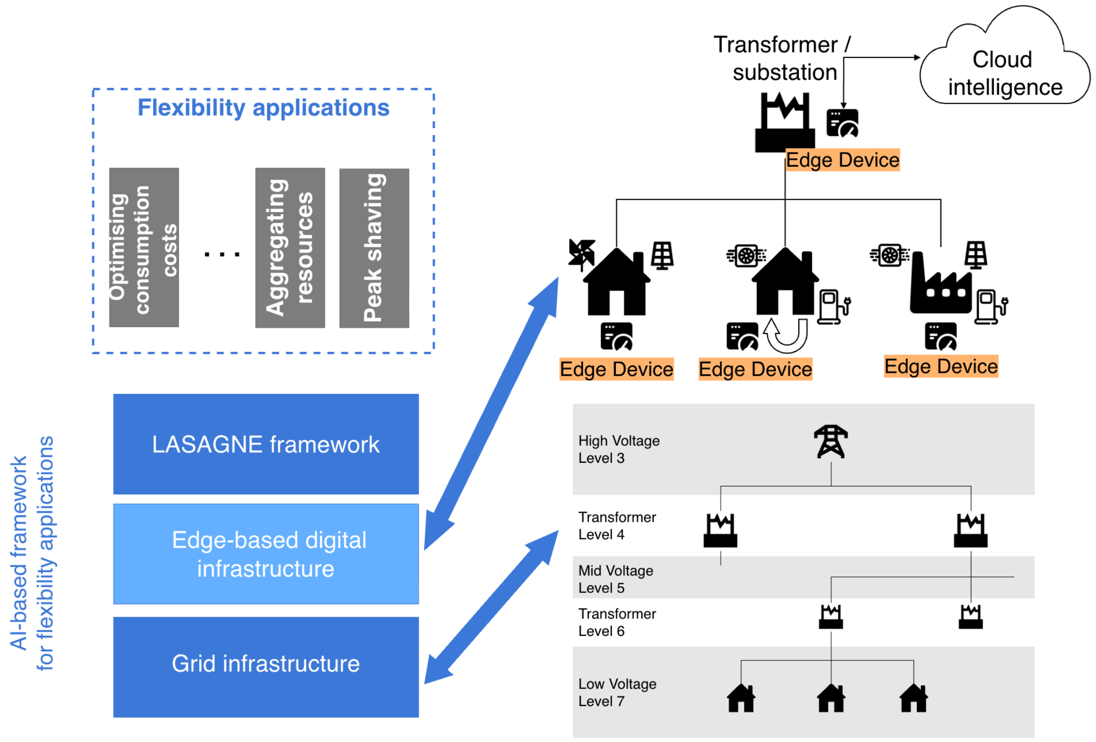

Fig. 1.1: The LASAGNE eco-system

The LASAGNE eco-system of Fig. 1.1 is composed of:

1. A Grid infrastructure: A microgrid (Local Energy Community:LEC) is deployed within a grid cell and spans mid- to low-level voltages (level 5 through 7). A microgrid is, in general, composed of households and factories employing several power sources (solar, wind, geothermal) and appliances (heat pumps, EV charging stations, HVAC, etc.).

2. An edge-based digital infrastructure: Low-GED (L-GED) and Mid-GED (M-GED) are able respectively to act on behalf of households and microgrids. L-GEDs and M-GED learn and anticipate the consumption/production of electric power at low (household) and mid (microgrid) levels and are then able to trade within the network of GEDs for energy exchange. L-GED and M-GED rely on the devices produced and marketed by CLEMAP. As described in the initial proposal, LASAGNE used the open source Nuvla/NuvlaEdge solution developed by SixSq[2] to manage the GEDs and deploy applications on these devices. The Nuvla platform is a secure, open edge-to-cloud management platform and marketplace, offering a complete operational solution that leverages edge and cloud computing with effective data management. The NuvlaEdge software manages and orchestrates edge computing resources through a container-based runtime environment built on the pre-installed Docker engine. It helps organisations deploy and manage applications at the edge of the network efficiently.

---

[2] https://sixsq.com/

3. LASAGNE Framework: Self-adaptive and context-aware consumption/production AI based forecasting algorithms. This framework contains a coordination model dedicated to the communication between the GED's.

The LASAGNE framework is the building block used to develop and deploy self-adaptive and context-aware forecasting algorithms that will themselves be used by the flexibility applications (grey boxes in Fig.1.1).The performance of the forecasting algorithm is continuously monitored: each time it is unable to take a reliable decision, it sends the data it failed to process to the cloud. The cloud then collects these datasets from different edge devices and triggers a new learning cycle, resulting in a new forecasting model which will then be pushed back to the edge devices (Fig. 1.2).

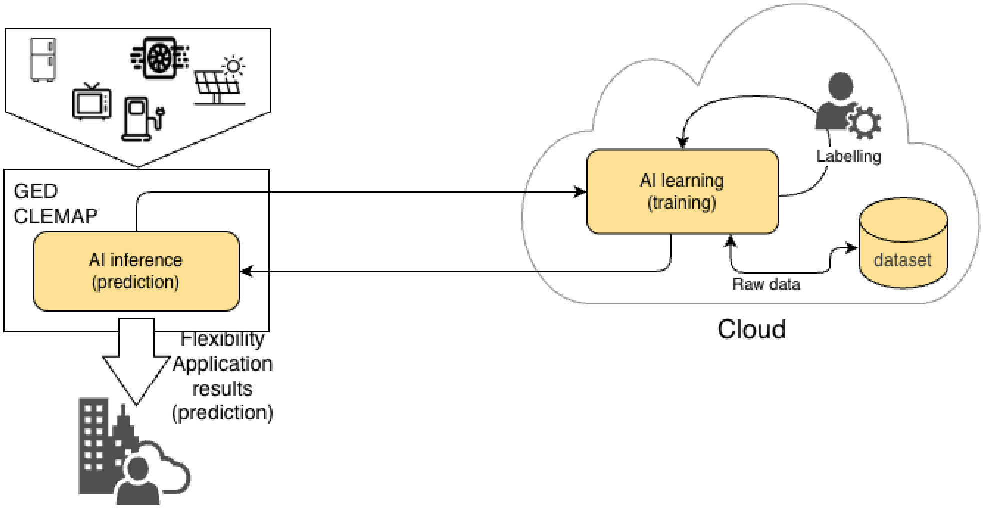

Fig. 1.2: Edge-to-Cloud approach to monitor and control forecasting energy algorithms deployed on the GEDs.

Another alternative is to enable peer-to-peer (Edge-to-Edge) communication among the GED to enable them to improve their intelligence. In this case:

1. The training takes place locally and the coordination model (part of the LASAGNE framework) is used as a communication backbone (Fig. 1.3).

2. A distributed federating learning approach is used to improve the forecasting performance. This work was the subject of a PhD thesis defended in 28th November 2025 by Mohamed Moussa at University of Geneva (UNIGE): *Towards decentralised machine learning predictions,for local energy communities*. The report of this PhD will be available soon on the website of UNIGE.

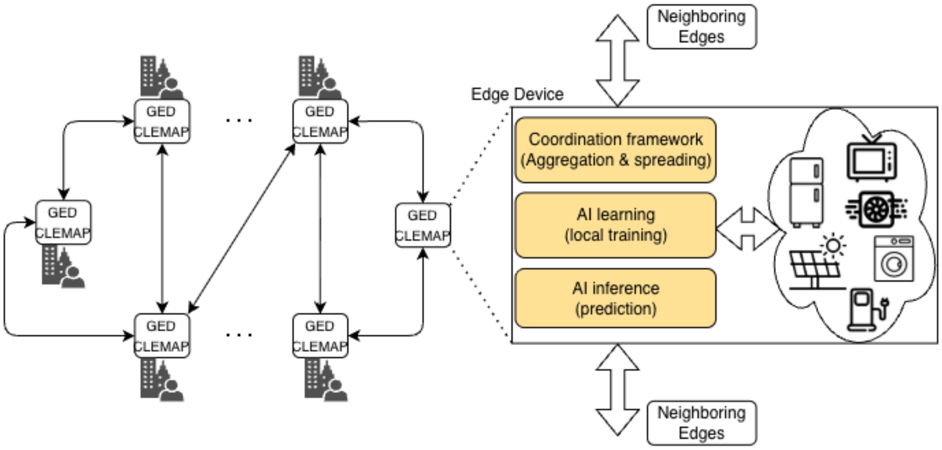

Fig. 1.3: Edge-to-edge approach to monitor and control forecasting energy algorithms

Although these two alternatives (Fig. 1.2 and 1.3) have not been fully implemented, they have served as a guiding principle throughout the LASAGNE project.

For the flexibility applications to spread out, they must be accepted and used by the different stakeholders. Whether these actors would tolerate such services (flexibility applications) is far from obvious and acceptance constitutes a potential barrier for the adoption and diffusion of our service. For this purpose, LASAGNE focuses on social acceptance of flexibility applications in the electricity sector. Because the design/development of flexibility applications and their deployment are far from standardised and many variants are conceivable, a series of realistic combinations were designed. GEDs are defined in various ways, depending on features such as direct load control, prediction and negotiation capacities. LECs are characterised by various features, such as how prices would be determined (by P2P markets or by a central authority) and the extent of the interaction between the power grid and the LEC. Each combination of features would yield different outcomes in terms of technical efficiency (electricity conservation, load shifting, etc.), but would also differ in terms of acceptability by the stakeholders.

This report is organised as follows: Chapter 2 describes how the CLEMAP device (GED) is customised, from hardware and software perspectives, to meet the requirements of LASAGNE. Chapter 3 details the different Machine Learning models tested within LASAGNE. The Long Short-Term Memory (LSTM)[3] model is selected to build the LASAGNE framework (Fig. 1.1). The chapter also discusses the different approaches used to execute Machine Learning algorithms on distributed data.

Chapter 4 describes an ML pipeline tool and a visualisation interface developed within LASAGNE. Users such as software developers can use the ML pipeline tool to experiment different forecasting configurations while end-users can use the visualisation interface to visualise forecasts and collected data. Chapter 5 describes the coordination model used to make the GEDs coordinate together to exchange information and improve their forecasting (Fig. 1.3).

---

[3] Long Short-Term Memory: https://www.bioinf.jku.at/publications/older/2604.pdf

Chapter 6 explores the concept of social acceptability, outlining a structured approach that transforms community insights into measurable preferences for the design of user-centric energy communities. Chapter 7 deals with the deployments in Sweden and Switzerland, and the marketplace concept developed within LASAGNE. Dissemination activities in 2024 and 2025 are detailed in chapter 8. Finally, chapter 9 presents the outlook for the LASAGNE project.

As a reminder, the milestones of the project, as detailed in the initial proposal, are:

1. M1: Validation of social attributes for flexibility applications
2. M2: Specification for edge-to-cloud implementation of distributed learning and coordination models
3. M3: Power and flexibility negotiation application developed
4. M4: Functioning lab version of Edge-to-Cloud marketplace platform according WP2 and WP3 specifications
5. M5: Deployed edge-to-cloud solution in Commune of Meyrin and ElectriCity

The table below details the links between the 9 chapters of this report and the project WPs and milestones.

Chapters	**WPs**	**Milestones**
Section 6 of the last progress report	WP4	M3
Chapters 2, 3, 4 & 5	WP3, WP4	M2
Chapter 6	WP2	M1
Chapter 7	WP5	M4, M5
Chapters 8 & 9	WP1	Dissemination and Knowledge Community Standard

Table 1.1: Correspondences between chapters WPs and milestones

Chapter 2: CLEMAP Device Customisation

Authors: HES-SO//HEPIA (Raoul Dupuis, John White, Nabil Abdennadher), CLEMAP (Gino Agbomemewa)

2.1 Introduction

This chapter examines the updates that need to be made to the current version of the CLEMAP device in order for it to meet the requirements of the LASAGNE project.

Section 2.2 provides a technical description of the CLEMAP device. Section 2.3 details the components of the LASAGNE framework to deploy on the same device. Section 2.4 summarises an experimental study to evaluate the performance of the current version of CLEMAP (based on Raspberry Pi 3) when deploying the edge-part of the LASAGNE framework (minimum requirement). This study shows that the Raspberry Pi3 version of the CLEMAP device cannot support the functionalities of the LASAGNE framework as described in section 2.3. Finally, section 2.5 details the software and hardware improvements carried out on the CLEMAP device to support the edge part of the LASAGNE framework.

2.2 The CLEMAP device

Initially, the CLEMAP device is composed of a custom three-phase measurement PCB and a Raspberry Pi 3. The electronic board, developed by CLEMAP AG and produced by HemarGroup AG in Mendrisio, Switzerland, connects to non-invasive current sensors (split-core transformers respectively Rogowski coils) and to the three phase conductors plus neutral in the electrical cabinet, where it is mounted on a DIN rail to measure per-phase voltage and current and derive active, reactive and apparent power as well as the power factor and phase angle (Fig. 2.1). The Raspberry Pi 3 is attached via its GPIO interface to this board and acts as the grid-edge controller: it samples the measurements at about 12 Hz, runs CLEMAP’s embedded software, and periodically publishes the data in a compact binary format over MQTT on two topics (with messages every 10 seconds and every minute). Through either Ethernet or Wi-Fi, the Raspberry Pi then forwards these measurements to the CLEMAP cloud backend, where they are visualized in the CLEMAP Floem web portal and can be combined with data from other CLEMAP devices for energy monitoring, analysis and future control or flexibility services in smart-grid and micro-grid scenarios.

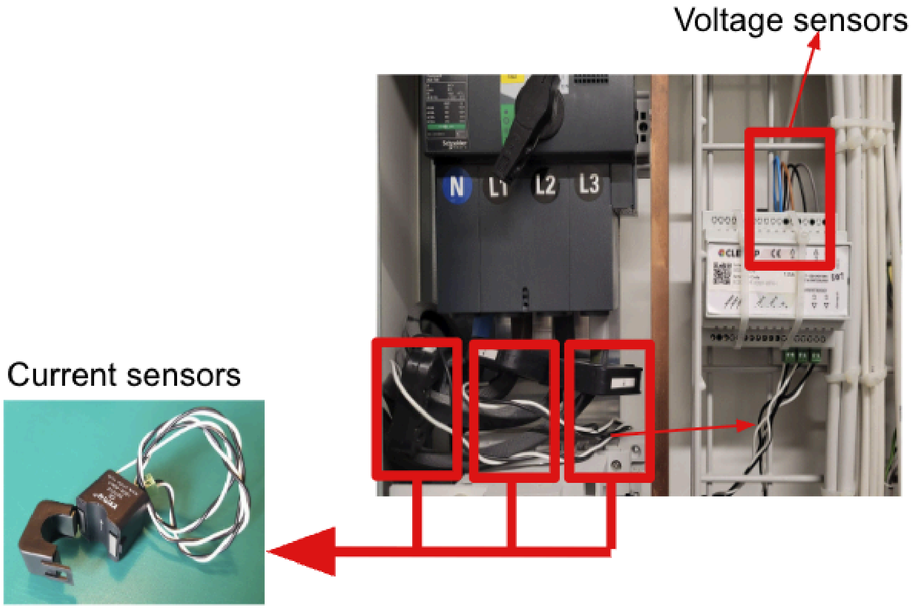

Fig. 2.1 : Typical cabinet installation of a CLEMAP device (To complete)

2.3 The edge components of the LASAGNE platform

The Fig. 2.2 outlines the main functional components deployed at the CLEMAP device and their interactions with the flexibility applications. The system integrates six key modules:

1. A digital Twin and Coordination framework,
2. A forecasting algorithm,
3. A local database (DB),
4. A data-gathering module,
5. A SmartGrid-Ready interface library (enabling interaction with home appliances and distributed energy resources)
6. NuvlaEdge (to manage and orchestrate the previous modules).

Together, these components support the implementation and deployment of flexibility applications for residential and small-scale energy systems (represented by "Flexibility Applications" in Fig. 2.2). The only module which was not deployed

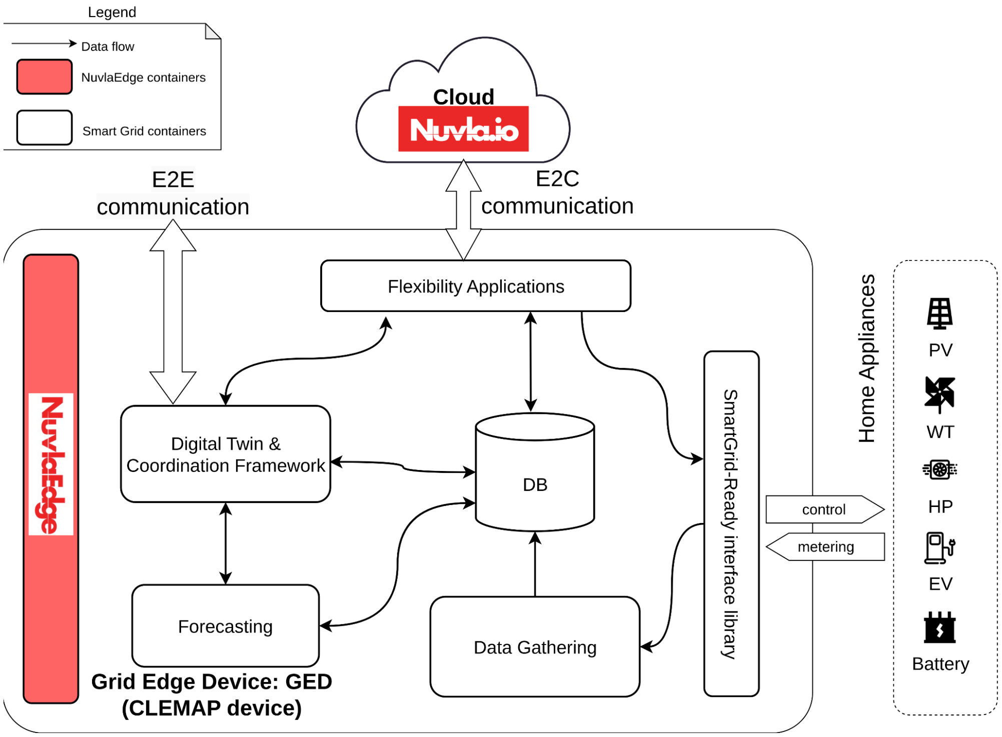

Fig. 2.2 Architecture of Grid Edge Device

SmartGrid-Ready Interface Library

The SmartGrid-Ready interface library enables seamless communication between the CLEMAP device and various home appliances such as photovoltaic systems, wind turbines, heat pumps, electric vehicles, and batteries. It supports both metering (e.g., power, energy, and state-of-charge retrieval) and control actions (e.g., setpoints, on/off commands, or operational mode adjustments). The library abstracts technology-specific communication protocols, like Modbus or proprietary EV charger protocol, providing a unified interface used by the data-gathering module and flexibility applications.

Data Gathering Module

The data-gathering component is responsible for collecting measurements from home energy devices and sensors through the SmartGrid-Ready interface library. It normalises, timestamps, and stores this information in the DB, as represented in Fig. 2.3. It also provides processed inputs to the forecasting engine and the Digital Twin. By ensuring reliable, low-latency acquisition of metering and status information, this module guarantees that the coordination platform operates on accurate and up-to-date data.

Database (DB)

The database stores all relevant operational data required by the GED ecosystem. This includes historical and real-time measurements, forecasting outputs, model states of the Digital Twin, and application-level data exchanged with cloud-based services. The DB ensures data persistence, efficient retrieval, and structured access for all components, making it a central element for analytics, decision-making, and coordination.

Forecasting module

The forecasting module predicts the energy consumption of the assets and appliances connected to the edge device. As an example, the graph in Fig. 2.2 illustrates a prediction model (XGboost) executed locally on the CLEMAP device, which forecasts the next 24 values with a measurement interval of 10 seconds. The blue curve represents the measured values, the red curve shows the most recent predictions, and the green curve corresponds to the recorded history of past predictions. The xgboost algorithm is not as precise as LSTM, but it is lighter to execute.

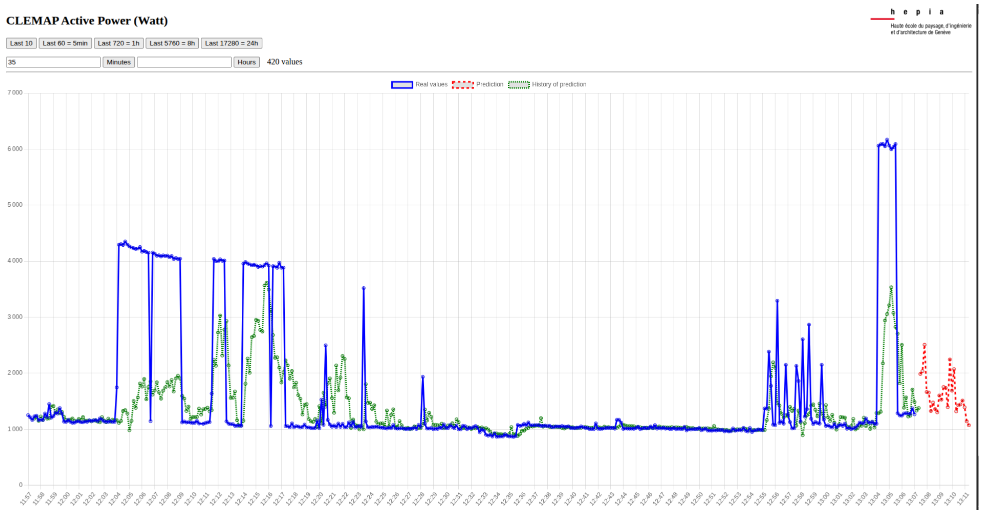

CLEMAP Active Power (Watt)

35 Minutes

420 values

Fig. 2.3 : Demo of real time data of actual and forecast (XGBoost) power measurement running on the CLEMAP device.

Digital Twin & Coordination framework

This module is the backbone that enables a network of GEDs (CLEMAP devices) to exchange information and improve their forecasting performance through a distributed federated learning and an edge-to-edge approach. It is assumed to be the implementation of the Edge-to-edge approach represented in Fig. 1.3. Digital twins (DTs), attached to the coordination platform, work as software agents coordinating their activities through the coordination platform. DTs work on behalf of GEDs (as consumers, producers of energy or as forecasting agents). As shown in Fig. 2.2., thanks to the coordination platform they can communicate and coordinate their actions with DTS at remote GEDs (E2E Communication); process data stored in the DB; and exchange information with the Applications.

NuvlaEdge

The NuvlaEdge software aims to provide a platform for managing and orchestrating edge computing resources. It helps organizations deploy and manage applications at the edge of the network efficiently. In LASAGNE, NuvlaEdge was used to monitor and control the deployment of the six basic modules in Fig. 2.1. NuvlaEdge assumes that the modules deployed on the edge are containerised.

The architecture described in Fig. 2.2 enabled us to implement and deploy the edge-to-cloud and edge-to-edge solutions presented in Fig. 1.2 and 1.3.

2.4 Deployment of a minimal version of LASAGNE framework on the CLEMAP device

This section evaluates the performance of the Raspberry Pi 3 based CLEMAP device, when running containerised energy-related services and, optionally, the NuvlaEdge module. Metrics such as CPU, memory, processes, temperature and container resource usage are collected using Prometheus[4], node-exporter[5] and Google cAdvisor[6].

Four software modules (Fig. 2.2) are part of this experiment (the two NuvlaEdge modules are part of the red NuvlaEdge box in Fig.2.2):

- The CLEMAP Edge Firmware (Smart Grid-Ready Interface library and Data Gathering in Fig. 2.2) retrieves measured values of voltage and current, computes power in alternating current, and stores them in the local SQLite Database.
- The Energy Forecasting, based on the LSTM model, predicts the next 10-seconds value using the last 4 values, each 10 seconds whenever a new value is collected. This is just a basic test that could be further changed based on the flexibility application requirements.
- NuvlaEdge_2.16.1 VPN Client (nuvlaedge-vpn-client): The nuvlaedge-vpn-client is responsible for picking up an OpenVPN configuration set by the nuvlaedge-agent and establishing a secure VPN tunnel with the VPN server, ensuring that the NuvlaEdge is always remotely accessible.
- NuvlaEdge_2.16.1 Agent (nuvlaedge-agent): The NuvlaEdge software allows the remote management of edge devices and deployment of containerised applications from the cloud Platform-as-a-Service Nuvla.io. The nuvlaedge-agent container is key in this process, as it is responsible for the NuvlaEdge activation, commissioning and monitoring procedures. All outgoing communication towards Nuvla passes through the agent. It notably sends heartbeat and telemetry reports on a regular basis with operational metrics from the

---

[4] https://prometheus.io/
[5] https://github.com/prometheus/node_exporter
[6] https://github.com/google/cadvisor

device (CPU, RAM usages, running containers). From Nuvla it receives requests with the actions users want to invoke on the edge (e.g. containerised application deployment, NuvlaEdge upgrade, etc.). The containerised applications management actions are done on the local Docker or Kubernetes container orchestrator engines.

The NuvlaEdge_2.16.1 Agent is the only required module of NuvlaEdge and constitutes its minimum configuration. The NuvlaEdge_2.16.1 VPN Client is optional; however, we include it in our testing as it was relevant in our deployment scenario.

Two configurations are tested:

- Deployment of the CLEMAP Edge Firmware and the Energy Forecasting on the CLEMAP device,
- Deployment of the four modules listed above on the CLEMAP device.

Each configuration was executed on a freshly rebooted RPI3B+ for two hours.

In configuration 1, CPU usage stabilises after initial spikes, with both containers requiring roughly 3–5% of CPU during normal operation. Memory consumption is significantly higher: the Energy Forecasting container uses about 248 MiB, representing ~27.5% of the available RAM, due mainly to ML model weight loading. Over time, memory usage gradually increases and the system heavily relies on swap, reaching nearly full swap usage during the test (a clear sign of memory saturation).

In configuration 2, introducing NuvlaEdge considerably increases CPU load. The VPN client alone consistently consumes around 20% CPU to maintain a secure connection. When deployments occur, CPU usage spikes significantly. The NuvlaEdge agent periodically shows abnormal high CPU demand, suggesting non-optimal behavior.

Memory usage again rises during deployment and runtime, driven especially by the Energy Forecast container. The combined pressure of all services increases swap usage and risks overall system instability.

Recommendations

The Raspberry Pi 3B+ is **heavily constrained**, mainly in memory. Even with a small set of services, RAM is nearly exhausted, forcing the system to rely on swap, which degrades performance and threatens long-term SD-card lifespan. CPU load is amplified significantly by NuvlaEdge operations.

The Energy Forecasting container, currently using a very simple model, already consumes large memory resources. Future model complexity would worsen the situation.

The recommended actions are:

1. Reduce memory usage of containers where possible (though forecasting optimisation may be limited).
2. Upgrade the hardware to a Raspberry Pi model with more RAM.
3. Avoid depending on swap, as increasing swap space will not mitigate performance bottlenecks.
4. Be aware that software evolution may soon exceed the capabilities of the current hardware.

2.5 CLEMAP hardware and software improvement

The CLEMAP device consists of two main boards: a processing board based on a Raspberry Pi and an interface board that manages sensors, power supply, and communication buses. These two boards are interconnected through the GPIO ports.

Hardware update

According to the experimental results of the previous section, the LASAGNE consortium decided to migrate to Raspberry Pi 4 or 5. The form factor of the Raspberry Pi 3, 4, and 5 remains identical, ensuring physical compatibility and seamless GPIO connections between the two boards, which means no mechanical changes are required when switching from one version to another. However, when moving from the Raspberry Pi 3 to the Raspberry Pi 4, it was necessary to modify the position of the Ethernet connector on the enclosure. On the Raspberry Pi 4 and 5, it is also possible to replace the SD card with a USB flash drive as the primary storage. However, this modification requires either replacing or physically altering the plastic enclosure to allow access to the USB port.

Software update

Following the hardware replacement, the software must be adapted to ensure full compatibility. This process begins with the use of the latest version of the Operating System (OS), Raspbian bullseye. Bookworm

Raspberry Pi 4

On the Raspberry Pi 4, the CLEMAP firmware was tested in a containerized environment and worked without modification. For the non-containerized version, only minor adjustments were required to achieve compatibility.

In addition, modifications were made to improve WLAN connectivity and to adapt the SPI driver for the measurement module located on the underlying board. Since the way Raspbian uses the *interfaces* module has changed, part of the firmware had to be rewritten to ensure proper integration.

Raspberry Pi 5

On the Raspberry Pi 5, the containerised CLEMAP firmware did not run successfully. Several modifications allowed the code to reach partial functionality. These changes mainly concerned the libraries responsible for GPIO interactions and for communication with the current measurement module. Additional adjustments were also required for OS-level interactions, including Wi-Fi configuration.

2.5 Conclusion

Experiments show that the Raspberry Pi 3 version of the CLEMAP device cannot support the two base components of the project - acquisition of electricity consumption measurements and Forecasting, when deployed with the NuvlaEdge module. The CLEMAP device had 1GB of RAM. The operating system and the CLEMAP measuring software consumed approximately 200 MB of the available 1 GB RAM. Further installing the NuvlaEdge agent and any extra containers running the base services caused the RAM usage to increase to such a level that the operating system started to use swap space. On such a small device with an SD card as the file system medium, this eventually caused crashes or freezing of the device. Automatic cleaning of swap space was not possible as this could lead to a crash of essential services.

This observation led to the conclusion that more RAM would be beneficial for the long-term operation of these devices when extra services are required or desired. Therefore, the consortium decided to migrate to Raspberry Pi 4 or 5. Because the effort required to achieve full compatibility with the Raspberry Pi 5 was too significant to complete within a reasonable timeframe, it was decided to standardize future deployments on the Raspberry Pi 4 equipped with 8 GB of RAM.

Chapter 3: Forecasting algorithms for electrical energy consumption

Authors : HES-SO//HEPIA (Chételat Jérôme, John White, Nabil Abdennadher), CLEMAP (Fabian Weiersmüller)

This chapter describes the Machine Learning (ML) models and Federated Learning approaches tested as part of the LASAGNE project. Section 3.1 presents the ML models studied for predicting electricity consumption at the building and household levels, while Section 3.2 presents an ML algorithm for predicting consumption at charging stations. Finally, section 3.3 details four approaches used to aggregate distributed data coming from a set of GEDs and use them during the training process. The idea here is to implement an approach that takes into account data distribution & privacy, and ensures “acceptable” prediction accuracy for the intended flexibility applications.

3.1 Forecasting algorithms for domestic energy consumption

This section details a set of ML algorithms used for time series forecasting. To study the effectiveness of these algorithms, it is necessary to have data covering at least one year. However, when the project was launched, the deployment data described in Chapter 7 was not yet available. We therefore decided to use similar data provided by a Swiss energy provider: data collected from more than 1 '000 smart meters during a period of two years. The acquisition frequency is one measurement every 15 minutes. It should be noted that the acquisition frequency used in our deployments is one measurement every 10 seconds.

The study conducted in this section can therefore be easily transposed using the data collected in LASAGNE deployments.

Five forecasting models were tested: XGBoost, LSTM, DeepAR, Temporal Fusion Transformer (TFT), and Prophet. This section discusses their conceptual foundations, their operational characteristics, and their suitability for microgrid and household-level forecasting tasks using real data from smart meters. The models were selected because they represent the most widely applied classes of forecasting algorithms in modern energy analytics: tree-based gradient boosting, recurrent neural networks, probabilistic generative models, transformer-based sequence architectures, and statistical forecasting methods.

Across these models, we evaluated prediction accuracy, complexity, and compatibility with federated learning paradigms. The broader objective is not only to compare their forecasting performance but also to identify which algorithmic family offers the best trade-off between accuracy, computational feasibility, and deployability on Grid Edge Devices (GEDs), which have constrained memory and processing capacity.

The results of these evaluations appear in Table 3.1 (RMSE and MAE scores) where LSTM emerges as the most suitable solution for subsequent federated learning research, outperforming or matching other models in accuracy.

3.1.1 eXtreme Gradient Boosting (XGBoost)

XGBoost is an ensemble-learning method that constructs additive trees to minimize prediction errors using gradient boosting. It is traditionally strong for tabular and structured data, offering excellent efficiency and fast inference. XGBoost has several advantages:

- good baseline performance in energy forecasting,
- low latency,
- high training efficiency.

However XGBoost is not natively compatible with federated learning. Integrating XGBoost into FL requires specialised secure aggregation mechanisms or custom adaptations, making it less practical for the context of LASAGNE.

3.1.2 Long Short-Term Memory (LSTM)

LSTM is the primary neural architecture for sequence modeling. Its ability to capture long-term temporal dependencies via controlled memory gates makes it particularly suitable for forecasting household-level consumption with daily, weekly, and seasonal patterns. Moreover, LSTM has natural compatibility with federated learning, since model parameters (weights) can be exchanged and aggregated across edge devices and strong performance with modest computational requirements.

3.1.3 DeepAR

DeepAR is a probabilistic modeling technique leveraging autoregressive recurrent networks. It is capable of generating full predictive distributions rather than point estimates, making it useful for uncertainty quantification. The two main drawbacks of DeeAR in the context of LASAGNE are (1) high computational cost unsuitable for GED-level deployment and (2) increased training complexity due to probabilistic modeling.

3.1.4 Temporal Fusion Transformer (TFT)

TFT represents a state-of-the-art transformer-based sequence model combining LSTM encoders with multi-head attention, gating mechanisms, and interpretable feature selection. While theoretically powerful, TFT suffers from pronounced operational limitations:

- extremely high memory and compute footprint,
- slow training speed,
- impracticality for edge deployment.

3.1.5 Prophet

Prophet is a decomposable time-series model that captures trend, seasonality, and holidays through additive components. Its strengths include interpretability and ease of use, but its main limitation is its lack of capacity to model complex temporal dependencies. While appropriate for business analytics with strong periodic structure, Prophet proves insufficient for fine-grained energy forecasting.

3.1.6 Comparative Evaluation

As shown in Table 3.1, we tested the forecasting accuracy in terms of RMSE and MAE. The scores were calculated for 100 smart meters (households) chosen randomly and the average of these 100 smart meters is reported as the final RMSE and MAE scores. The models were trained on data between January 2021 and April 2022 and evaluated on a test set for the time May to June 2022. All models were trained with 1 hour of lookback to predict the next hour's net active power (1-hour horizon).

Model	**RMSE (kW)**	**MAE (kW)**
LSTM	**0.381**	**0.209**
XGBoost	0.392	0.225
Transformer	0.474	0.337
DeepAR	0.533	0.429
Prophet	0.612	0.487

Table 3.1 : RMSE and MAE (expressed in kW) score results using 100 households (assumed to be GEDs), based on different models

This evaluation demonstrates the LSTM performs best overall. XGBoost is competitive but incompatible with FL. Finally, DeepAR, TFT, and Prophet underperform while requiring far greater computational resources.

3.2 Energy consumption forecasting for Electric Vehicle chargers

With the electrification of mobility the grid reaches its limits of capacity, in particular at the extremities of the grid. GEDs within this application help to reduce the load intensity while controlling chargers. This section outlines the application designed to forecast the power consumption of Electric Vehicle (EV) chargers. These chargers are connected to CLEMAP devices, which measure both the electrical current drawn from the grid and the operational status of each charger (e.g., charging, idle). Data collected over one year serves to train a machine learning model, enabling it to predict the energy usage of clustered EV charging stations, collectively referred to as a Virtual Power Plant (see Fig. 3.1).

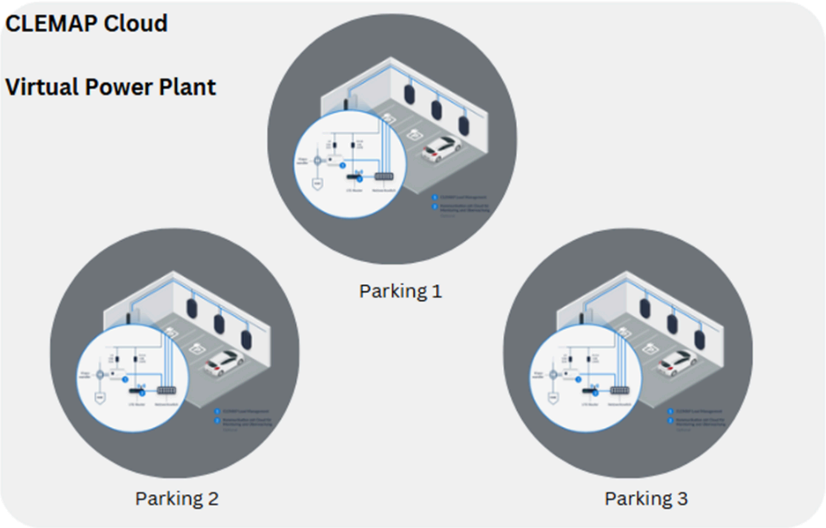

Fig. 3.1: Aggregating here in the example three parking with charging infrastructure to one Virtual Power Plant

3.2.1 Data and Preprocessing

The timeframe of the available data is between January 1, 2024 and July, 2025. Each record consists of a measurement that is done every 30 seconds on the three electric phases used by the EV charger, the timestamp at which the measurement took place and the status of the EV charger. Throughout the dataset, the volume of collected data points steadily increased as additional devices were deployed. In total, approximately 1.5 billion raw data records were collected and subsequently compressed into parquet format before preprocessing.

3.2.2 Data cleaning and preparation

Raw data requires preprocessing to eliminate unstable or redundant information before being used in machine learning.

Fig. 3.2 and Fig. 3.3 exhibit the significant portion of the data representing idle states.

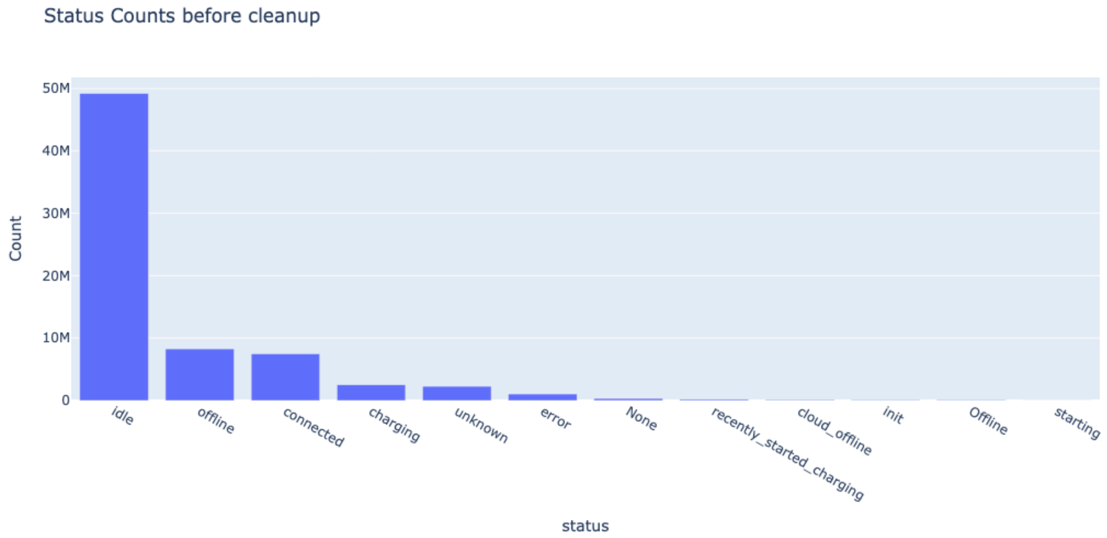

Fig. 3.2 : Counts of the statuses over the month of May 2024

Fig. 3.3: Device recording only periods of offline status or idle states.

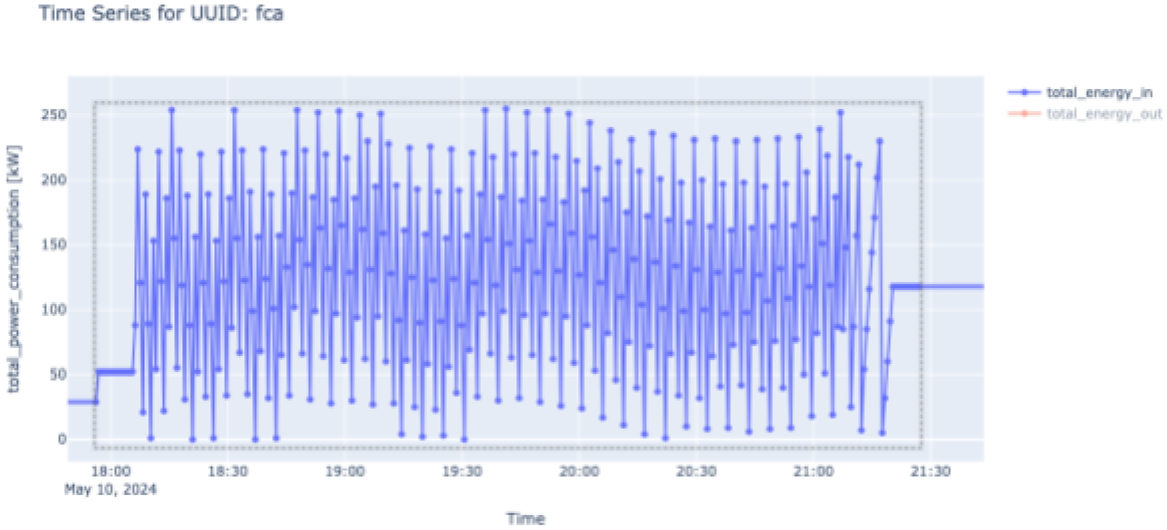

Fig. 3.4 : Jittery data example over a period of charging time.

To reduce dataset size and smooth noisy data as displayed in Fig. 3.4, the original dataset was downsampled into 15-minute intervals, resulting in approximately 25 million records. Special attention was necessary due to categorical status data collected by the CLEMAP devices. While numerical data could be effectively aggregated using these statistical measures, categorical data such as device statuses required an alternative strategy. Specifically, for categorical variables, either the first recorded status within each interval or the most frequently occurring status could be selected to represent the aggregated data point.

3.2.3 Feature engineering

The dataset contains current measurements across the three phases (L1, L2, L3). An additional column representing power consumption in watts was generated and normalised to optimize the model's performance when recognizing specific consumption curve signatures.

During the downsampling step, statistical metrics including median, standard deviation, first quartile (Q1), third quartile (Q3), minimum, and maximum were calculated individually for each of the three phases. These metrics were specifically computed to facilitate uncertainty estimation after the forecasting process.

Clustering techniques were employed to classify EV charging stations into groups based on similarities in their power consumption patterns. Stations were categorized into five distinct types: Home, Enterprise, Commercial, On-street, and Public.

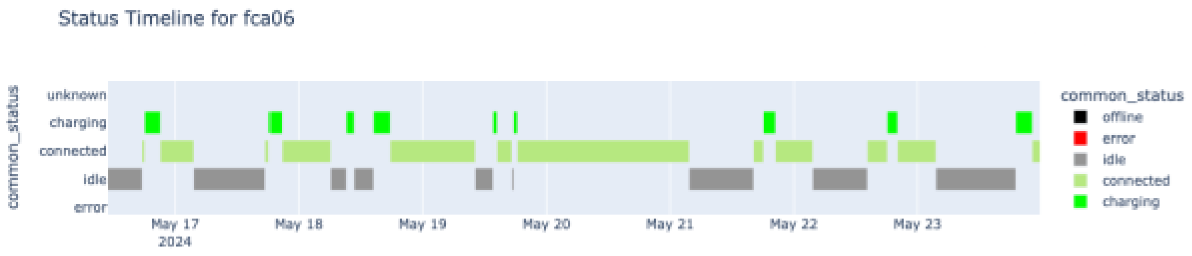

Fig. 3.5 : Example of an EV charging station that would be located in a residence

Fig. 3.5 illustrates a typical scenario of an individual with a regular job, unplugging their electric vehicle in the morning upon leaving home and reconnecting it in the evening to recharge. Additionally, the figure highlights a noticeable change in this routine during the weekend (May 19 and 20, 2024), demonstrating a different usage pattern.

Fig. 3.6 and 3.7 show EV chargers on public streets and enterprise respectively.

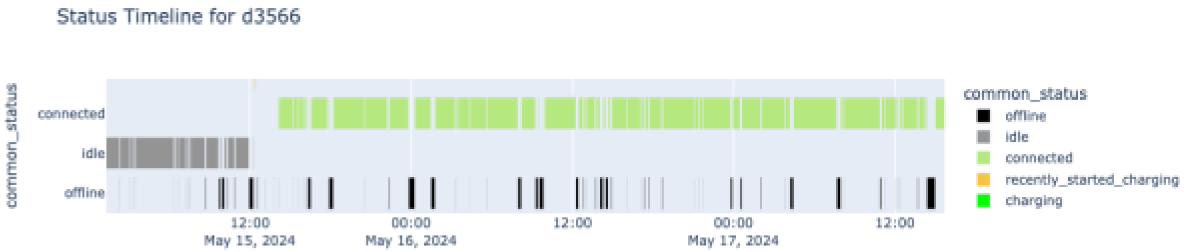

Fig. 3.6 : Example of an EV Charger that can be located on the street or in a parking lot.

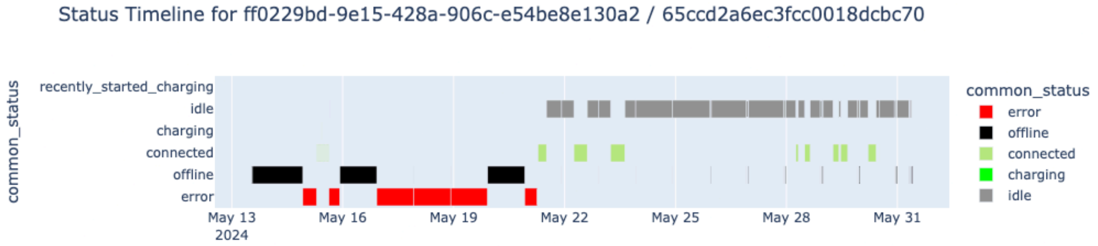

Fig. 3.7 : Example of an EV Charger that can be located at an enterprise.

This labelling of EV Chargers can improve the forecasting by predicting if a routine is in place or not depending on the type of charger. See Fig. 3.6 and Fig. 3.7 here above which show how the behavior in the car charging is different. The usage generates a load profile of the chargers which can be highly different dependant to location, see Fig. 3.8.

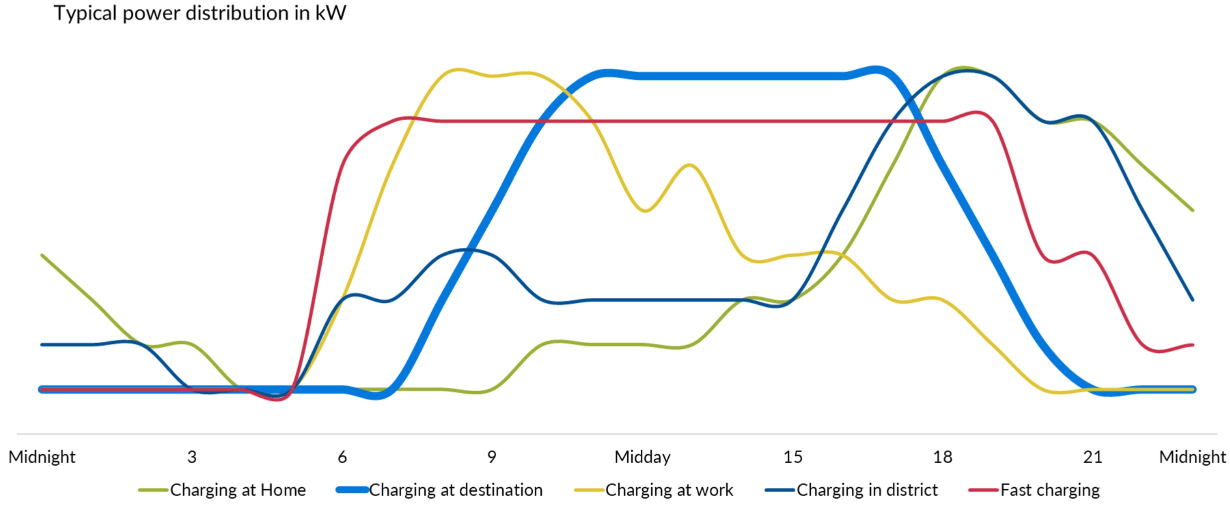

Fig. 3.8: Typical power distribution for different charging behaviors in different charging locations

3.2.4 LSTM Model

The implemented model uses PyTorch and consists of two LSTM layers, each containing 128 hidden layers. It predicts future values based on a lookback window of 7 days (equivalent to 672 data points in 15 minutes granularity), with a forecast horizon of 7 days ahead, capturing daily patterns within the time series data.

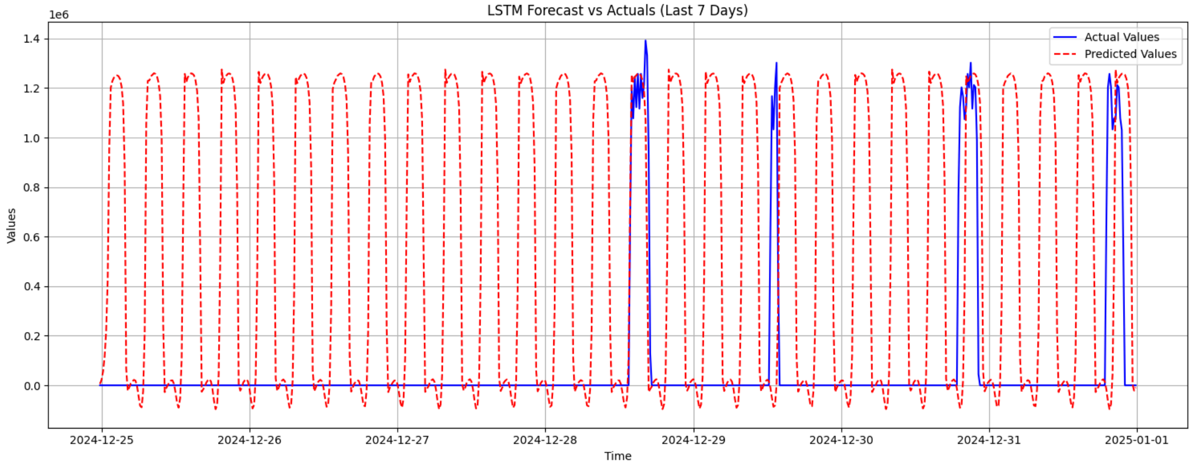

Fig. 3.9 : Prediction of the current model on the test dataset. The y-axis scale represents a power in a scaled format.

As shown in this forecast, the dotted line exhibits periodic oscillations but fails to closely match the curve of actual values. This discrepancy might be caused by the model becoming trapped in a local minimum, preventing it from finding a more optimal fit and generalisation of the data.

3.4 Federated Learning

This section is a summary of chapters 5 and 6 of the Doctoral Thesis of Mohamed Moussa (presented on 28th November 2025): Towards decentralised machine learning predictions for local energy communities. The report of this PhD will be available soon on the website of UNIGE. Some parts of this section have been copied from chapters 5 and 6 of the thesis report.

This section introduces and evaluates a Federated Learning (FL) approach designed to forecast household and microgrid electricity consumption while preserving user privacy. The work uses the same data used in Section 3.1.6 and investigates how FL improves accuracy, generalisation, and privacy preservation.

Traditional ML approach (Fig. 3.10), called here “multi-device approach”, requires aggregating all raw data in a central server. This approach is effective in the case of new GEDs (that do not have any data): the model can be applied to the new GEDs without any learning. However, it has three drawbacks:

- data privacy is not respected,
- communication between the GEDs and the server can be costly when the number of GEDs increases,
- aggregating data from GEDs that are not 'similar' can affect the performance of the predictive model.

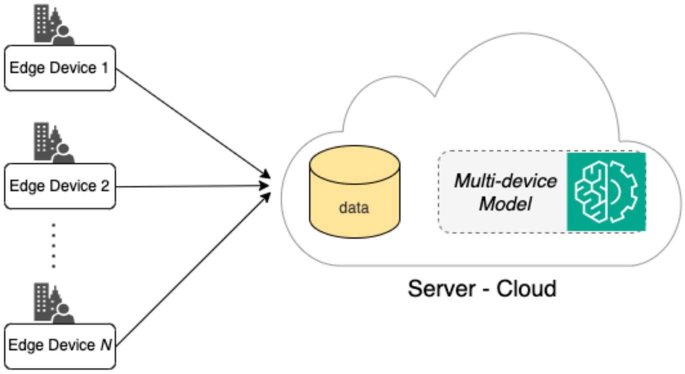

Fig. 3.10: Multi-device approach

The second tested approach is the single device scenario (Fig. 3.11). In this model, each GED trains its own personalised local model using only its own historical data. Here, predictive models are customised and tailored to the relevant GEDs. Privacy data is fully respected and specific anomalies/habits are easily detected. However, the model presents scalability issues and maintenance overhead and cannot be applied when the number of GEDs increases.

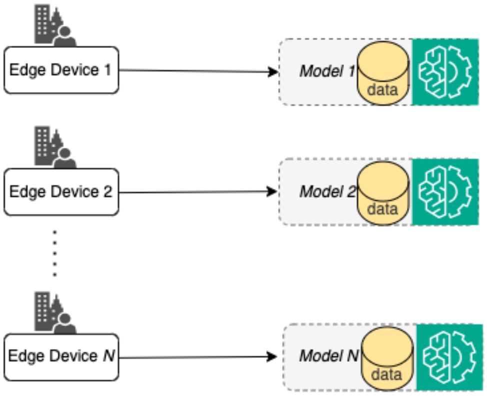

Fig. 3.11: single-device approach

To overcome the shortcomings of these two approaches, LASAGNE experimented with a Federated Learning approach. Here the GED trains locally and sends only model updates (weights or gradients) to a central server. The server aggregates updates and broadcasts back an improved global model. We called this approach: Centralised FL (CFL) since the server centralises the weights received from the GEDs. The only drawback to this approach is that aggregation is performed using all GEDs. As explained above, aggregation with non-similar GEDs does not make sense and can negatively impact the performance of the predictive model.

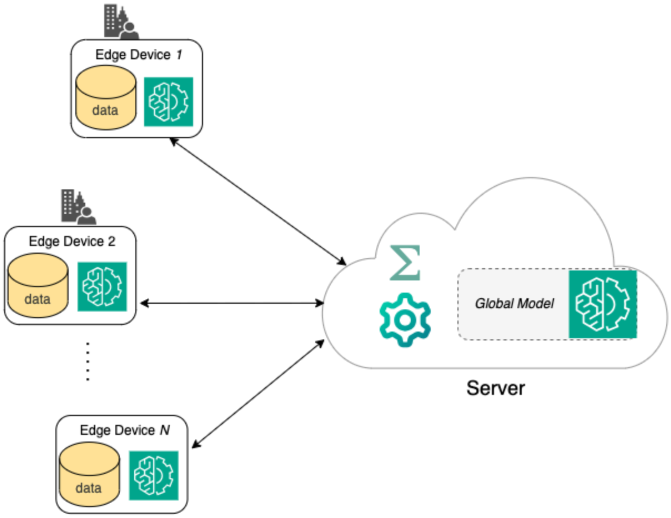

Fig. 3.12: Federated Learning approach

A final approach tested as part of the LASAGNE project is to define clusters of GEDs that are assumed to behave in the same way. To this end, each GED has a signature that identifies it. GEDs with 'similar' signatures are assigned to the same cluster. Furthermore, in this approach, the server in Fig. 3.12 is removed and replaced by peer-to-peer communication between GEDs belonging to the same cluster. This approach is called decentralised federated learning (DFL)

The signature is calculated by averaging the consumption values for each day of the week. For each of these days (Monday through Sunday), we compute the average consumption for the 96 intervals (one measure for each 15 minute interval = 4 * 24 = 96). This means calculating the average of the 96 consumption values recorded on each Monday, then separately on each Tuesday, Wednesday, ... and Sunday. Fig. 3.13 represents raw data of a given GED used to calculate the signature represented in Fig. 3.14.

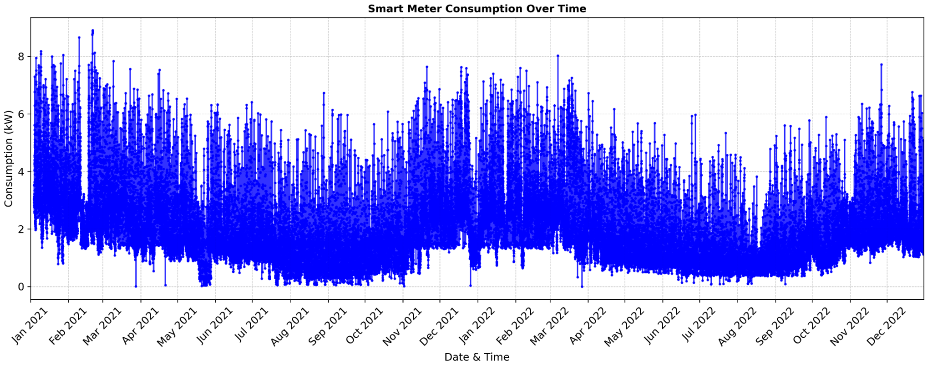

Fig. 3.13: Raw data of one smart meter (GED in our case).

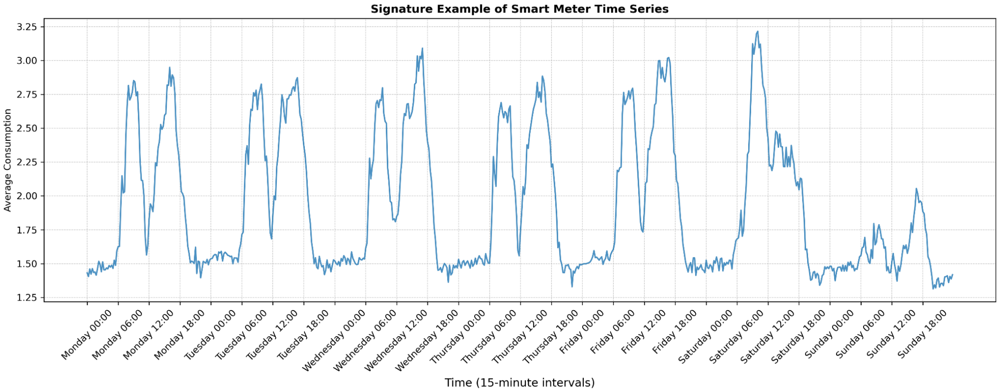

Fig. 3.14: Signature of the smart meter (GED in our case) represented in Fig. 3.13.

Each edge then calculates the distance between its consumption signature, and other nodes. It is obtained using a predetermined distance function that catches similarity between energy signatures. The distance calculated is then used to rank all possible communication partners. Finally, each edge selects the top *x* closest neighbors. *x* is a parameter that can control the communication overhead and the neighborhood size.

As specified in the previous sections, we use a 1-day lookback and a 1-day forecasting horizon to prepare the data for supervised learning. The training dataset spans from January 2021 to April 2022, and the test dataset covers May to June 2022. To assess the performance of each approach, we computed the mean of MSE and the standard deviation of all MSE scores for all GEDs (both expressed in kW). By using both mean of MSE and standard deviation of MSE scores, we can get a comprehensive understanding of the model's performance. The mean of MSE scores provides us with a summary of the overall accuracy (a representative value for the overall accuracy of the model in predicting the power load across all GEDs), while the standard deviation of MSE scores provides us with a measure of the consistency of the model's performance across different GEDs.

Table 3.1 shows that the multi-device approach performs better than the single-device approach. One possible explanation for the superior performance of the multi-device LSTM model is that it is able to capture the underlying patterns in the power load data more effectively than the single-device model. This may be due to the multi-device model being trained on a larger and more diverse dataset, allowing it to learn a broader range of patterns and features.

Number of GEDs	Mean MSE(± Std MSE) Multi-device approach	Mean MSE(± Std MSE) Single-device approach
50	0.239 (± 0.524)	0.436 (± 1.122)
100	0.319 (± 1.417)	0.651 (± 2.557)
200	0.273 (± 1.666)	0.678 (± 3.517)
500	0.417 (± 1.716)	0.648 (± 2.747)

Table 3.1: Multi-device vs. single device (MSE is expressed in kW)

Table 3.2 shows that the CFL approach achieves comparable performance to the multi-device approach. We can notice a slight improvement in CFL performance with the increase in the number of clients (from 5 to 50 GEDs), approaching the accuracy of the multi-device approach. CFL does not outperform the multi-device approach, but consistently outperforms single-device.

Approach	Mean MSE(± Std MSE)
Federated (5 clients)	0.105 (± 0.218)
Federated (10 clients)	0.104 (± 0.210)
Federated (20 clients)	0.103 (± 0.208)
Federated (30 clients)	0.103 (± 0.207)
Federated (50 clients)	0.103 (± 0.206)
Multi-device	0.103 (± 0.211)
Single-device	0.257 (± 0.577)

Table 3.2: Mean MSE and Std MSE (in kW) with the CFL approach.

In Table 3.3, the DFL variant with neighbor-based communication is clearly effective in improving the performance of the “all-to-all” approach. The all-to-all approach is the DFL where all GEDs exchange information with all GEDs. By communicating with just the most relevant neighboring nodes rather than all nodes in the network, DFL reduces communication overhead while improving prediction accuracy.

# Devices	**Multi-device**	**DFL (All-to-All)**	**DFL with x Neighbors**			
			5	**10**	**20**	**30**
50	0.040	0.065	0.035	0.034	—	—
100	0.045	0.075	0.035	—	0.034	—
200	0.056	0.131	0.034	—	—	0.033

Table 3.3: NRMSE for multi-device and Decentralised FL Approaches (1-hour Ahead Prediction - LSTM Model)

In Table 3.3, we did not test all the possible configurations for each set of devices, as we observed that there is no significant improvement when varying the number of performances. As an example, in Tables 3.3, some cells are empty (—), indicating configurations which have not been tested.

3.4 Conclusion

The experiments conducted in this chapter have shown that the LSTM model is best suited for LASAGNE (section 3.1). This model is also used for EV charging station forecasting in section 3.2. Section 3.3 shows that the Decentralised Federated Learning (DFL) approach can provide a solution to the problem of distributed and private data.

While exploring alternative modeling approaches—including more advanced architectures such as transformer-based models—remains of interest, the immediate priority is to ensure transparent and quantifiable accuracy. Once improved results are achieved and validated, the statistical metrics will serve to clearly communicate prediction uncertainty to end users. Any potential integration into CLEMAP's FLOEM [7]platform including commercialization with a licensing model should be considered exploratory at this stage.

---

[7] FLOEM is a data management platform that provides access to CLEMAP devices data.

Chapter 4: Configurable ML Pipeline for Electricity Consumption Forecasting

Authors: HEPIA (Mohsen Salimi, John White, Raoul Dupuis), KTH (Carlo Fischione)

4.1 Introduction

This chapter outlines a pipeline developed for training machine learning models to forecast electrical consumption. The ML pipeline is assumed to support any kind of dataset. The ML pipeline tool was concretely experimented with a dataset collected from CLEMAP devices deployed in Meyrin and Chêne-Bougeries (see details in Chapter 7). Data is accessed through the FLOEM platform[8] which manages data streams from CLEMAP devices.

4.2 ML pipeline for energy forecasting

The whole system is designed as a multi-stage pipeline, with each stage performing a distinct transformation on the data. This modular architecture promotes maintainability, enables independent optimisation of each stage, and facilitates experimentation with alternative approaches. Fig. 4.1 illustrates the overall architecture of the four-stage pipeline along with the orchestration and tracking infrastructure. The pipeline consists of four primary stages: data cleaning, feature engineering, data preparation, and model training.

Pipeline orchestration is achieved through Data Version Control (DVC)[9], a version control system specifically designed for machine learning projects. DVC tracks both code and data dependencies, ensuring that any change to parameters, code, or input data triggers the appropriate downstream recomputation. This approach guarantees reproducibility.

Experiment tracking and model performance monitoring are implemented using MLflow[10], an open-source platform for managing the machine learning lifecycle. MLflow automatically logs hyperparameters, performance metrics, and training artifacts, enabling systematic comparison across experimental configurations.

---

[8] floem.clemap.com
[9] https://dvc.org/
[10] https://mlflow.org/

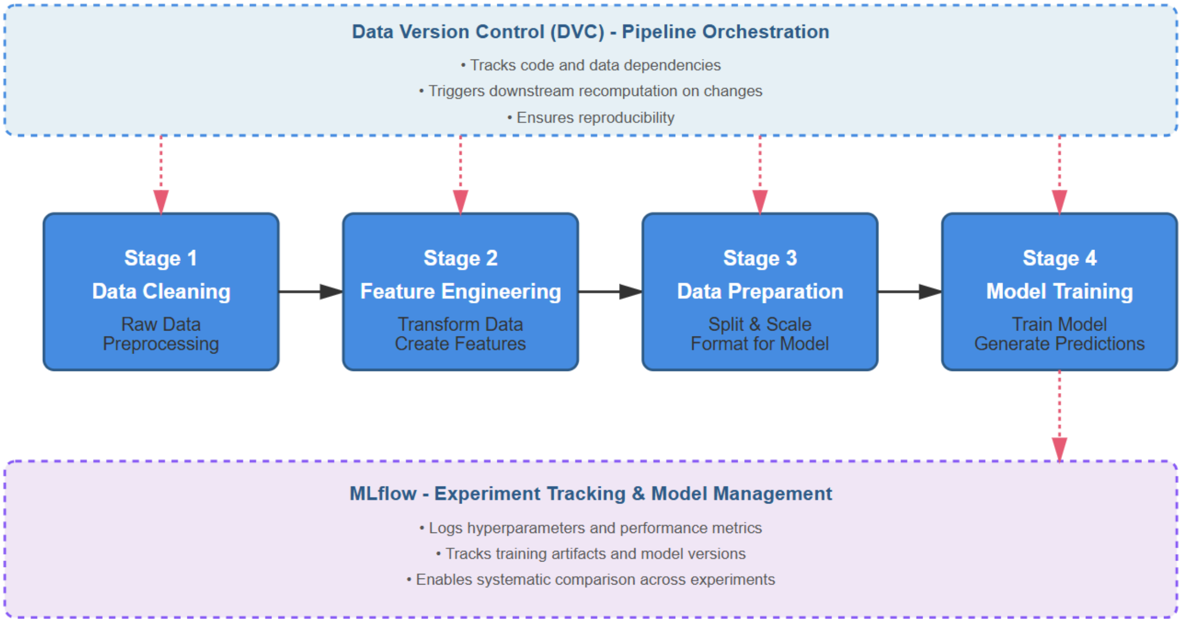

Fig. 4.1: Overview of the four-stage machine learning pipeline for electricity consumption forecasting.

- **Data Cleaning**
  This subsection outlines the procedures used to clean and validate the raw data obtained from the FLOEM platform. It describes how missing values, duplicated entries, irregular timestamps, and outliers are detected and handled to ensure the reliability of the input dataset.

- **Feature Engineering**
  Here, the focus is on constructing meaningful features from the cleaned time-series data. The subsection details the extraction of temporal indicators, lagged variables, statistical aggregations, and any domain-specific transformations that enhance the predictive capacity of the models.

- **Data Preparation**
  This part explains how the engineered dataset is transformed into a format suitable for machine learning. It covers normalization or standardization operations, dataset partitioning into training/validation/testing sets, and the generation of sliding windows for supervised forecasting tasks.

- **Model Training**
  The final subsection describes the model development process. It includes the training of forecasting algorithms, hyperparameter optimisation, evaluation procedures, and the use

of DVC and MLflow to orchestrate the pipeline and track experiments in a reproducible manner.

Data Cleaning

FLOEM Platform

The data used to train the forecasting model in this study originated from CLEMAP devices manufactured by CLEMAP. These devices capture comprehensive electrical parameters across three phases at multiple temporal resolutions. Access to this data was facilitated through the FLOEM platform, a data management platform that provides access to CLEMAP device data. The CLEMAP device captures six primary electrical parameters for each of the three phases:

- Active power (P) in watts
- Reactive power (Q) in volt-amperes reactive
- Apparent power (S) in volt-amperes
- Voltage (U) in volts
- Current (I) in amperes
- Power factor (PF) as a dimensionless ratio

For this study, data from 12 CLEMAP devices was collected. The deployment of these 12 Devices is detailed in chapter 7.

A challenge in working with the FLOEM platform was the limitation imposed on data retrieval operations. The platform's API restricts individual download requests to a maximum of 6 hours of high-resolution data. For comprehensive analysis requiring months or years of historical data, this constraint necessitated the development of a custom solution.

A Python-based custom data retrieval client was developed to address these limitations. This custom interface implements intelligent algorithms that automatically segment large time ranges into platform-compliant requests, handle authentication through JSON web Tokens, and ensure data integrity through validation checks. The client provides a seamless abstraction layer, allowing to request arbitrary time ranges without manually managing the underlying pagination and chunking logic.

Challenges with Raw Data

Raw data from CLEMAP devices' infrastructure invariably contains quality issues that, if unaddressed, can severely compromise model performance. The observed data quality challenges included:

- Incomplete days: Days with missing measurements due to communication failures or device maintenance
- Timestamp irregularities: Measurements not precisely aligned to expected 15-minute intervals
- Zero values: Erroneous zero readings indicating measurement or transmission failures

- Data gaps: Extended periods of missing data due to various operational issues

These quality issues necessitated implementing a comprehensive cleaning pipeline to ensure that only reliable, complete data was used for model training.

Cleaning strategy

A systematic cleaning approach was developed, consisting of several validation and correction steps.

First, timestamp validation was performed to identify and correct minor timing irregularities. Measurements within a configurable tolerance window (±2 minutes) of expected timestamps were snapped to the nearest standard interval, ensuring temporal consistency across the dataset.

Second, day-level completeness validation was implemented. A complete day at 15-minute resolution should contain exactly 96 measurements. However, to accommodate edge cases in which data loss occurred during the first or last hour of the day, the criterion was relaxed to accept days with at least 95 valid samples. Days failing to meet this threshold were excluded from the training dataset.

The cleaning process is implemented as a configurable Python module with comprehensive logging, allowing full transparency into the data quality assessment and cleaning decisions. Detailed reports documenting the number of records processed, rejected, and corrected are automatically generated for each cleaning operation.

For more information on the details of the cleaning methodology, please refer to the Git[11] repository of the project.

Feature Engineering

Effective machine learning for time series forecasting requires transforming raw temporal measurements into a rich feature space that captures relevant patterns and relationships. An extensive feature engineering pipeline was developed, generating over 100 features from the raw power measurements.

Temporal Features

Temporal features encode calendar and time-of-day information, enabling the model to learn periodic patterns associated with human activity cycles. Basic temporal components, including hour of day, day of week, day of year, week of year, month, quarter, and year, provide direct calendar information. Binary indicator features were created for weekends, specific days of the week, and time periods (working hours, morning peak, evening peak, night hours).

---

[11] https://gitedu.hesge.ch/lsds/phd

To better represent the cyclical nature of temporal phenomena, trigonometric encodings were applied to periodic features. Hour of day, day of week, month, and day of year were transformed using sine and cosine functions, creating continuous circular representations that properly capture periodicity.

Lag Features

Lag features provide the model with explicit access to historical values, capturing auto-regressive patterns in the time series. Lag features are created at multiple temporal scales: short-term (15 minutes, 30 minutes, 1 hour, 2 hours), medium-term (6 hours, 12 hours), and long-term (1 day, 1 week). This multi-scale lag structure enables the model to leverage both recent trends and longer-term periodic patterns.

Rolling Window Statistics

Rolling window features compute statistical aggregates over sliding temporal windows, capturing local trends and variability. For each power measurement, four statistics (mean, standard deviation, minimum, maximum) are computed over five different window sizes (1 hour, 3 hours, 6 hours, 12 hours, 1 day). These features provide the model with a summarized view of recent consumption patterns.

Rate of Change Features

To capture the dynamics of power consumption, several rate-of-change features are engineered: first and second differences, percentage change between consecutive measurements, velocity (change over 1 hour), and acceleration (rate of change of velocity). These derivative features help the model recognize rapid changes in consumption patterns.

Statistical Features

Beyond simple aggregations, several advanced statistical features are computed to characterize consumption patterns. These include the Hurst exponent to measure long-range dependence, sample entropy to quantify signal complexity, and approximate entropy to assess time series irregularity. Additional statistical descriptors include skewness, kurtosis, and the coefficient of variation computed over sliding windows.

Calendar Features

To capture the influence of socio-temporal factors, comprehensive calendar features are implemented. The pipeline automatically identifies public holidays specific to the canton of Geneva in Switzerland, including both fixed holidays (New Year's Day, National Day) and variable holidays (Easter, Ascension Day). Additionally, features marking school holidays and semester breaks are included, recognizing that educational calendar events significantly affect consumption patterns in university facilities.

The holiday detection system leverages the Swiss government's holiday database, ensuring accurate identification of both cantonal and federal holidays. These features allow the model to learn distinct consumption patterns associated with holiday periods versus regular operational days.

Feature Selection and Management

Given the extensive feature set, a configuration system was developed to enable selective feature utilization. Features are organized into categories (temporal, lag, rolling statistics, rate of change, statistical, calendar), and researchers can specify which categories to include in model training through simple configuration parameters. This modular approach facilitates systematic ablation studies to assess the relative importance of different feature types.

All feature engineering operations are implemented vectorised using pandas[12] and numpy[13], ensuring computational efficiency even for large datasets spanning multiple years.

Data Preparation

Dataset Structure and Splitting

Following feature engineering, the dataset is organised into a supervised learning format suitable for LSTM training. The fundamental unit of the dataset is a sequence-label pair, where the sequence (lookback window) contains historical features, and the label represents future power consumption values to be predicted.

The dataset was split into training, validation, and test sets using temporal partitioning to preserve the time series nature of the data. This temporal split ensures that the model is evaluated on genuinely unseen future data, providing realistic performance estimates. The standard split ratio of 70% training, 15% validation, and 15% test was employed.

Sequence Generation

The sequence generation process creates overlapping windows of historical data paired with corresponding future consumption targets. Two key parameters govern this process:

Lookback window: The number of historical time steps provided to the model as input context.

Prediction horizon: The number of future time steps to forecast.

The sequence generator creates samples by sliding the lookback window across the time series, with each sample representing an independent training instance. Sequences that span missing data or incomplete days are excluded from the dataset.

---

[12] https://pandas.pydata.org/
[13] https://numpy.org/

Normalisation

All features and target variables are normalised to improve training stability and convergence. Standardisation (z-score normalisation) is applied, transforming each feature to have a normal distribution (mean of 0 and a variance of 1). Critically, normalisation statistics are computed exclusively on the training set and then applied to validation and test sets, preventing data leakage.

For the inverse transformation of predictions back to original units, normalization parameters are stored and made available to the evaluation pipeline.

Model Architecture and Training

The forecasting model is based on a stacked LSTM architecture, a proven approach for sequence modeling tasks. The network architecture that has been tested consists of:

- Input layer: Accepts sequences of shape (lookback window, features)
- LSTM layer 1: 128 units with return sequences enabled
- Dropout layer 1: 20 dropout rate for regularization
- LSTM layer 2: 64 units without return sequences
- Dropout layer 2: 20 dropout rate
- Dense output layer: Linear activation with units equal to the prediction horizon

This architecture balances model capacity with computational efficiency. Developers can modify these parameters to observe the results of their changes. Stacked LSTM layers facilitate hierarchical feature learning. For example, when two LSTM layers are stacked, the first layer analyzes the sequence to identify simple time patterns, like subtle changes over time. The output from this initial LSTM layer is then provided to the second LSTM layer, which identifies more complex, higher-level patterns derived from the first layer’s findings.

Training Configuration

Model training employed the Adam optimiser[14] with an initial learning rate of 0.001. Mean squared error (MSE) was used as the loss function, appropriate for regression tasks. Training is conducted for up to 100 epochs, with early stopping based on validation loss.

Hyperparameter Management

All hyperparameters (network architecture, training configuration, feature selection, data processing parameters) are managed through a centralised YAML configuration file. This

---

[14] https://arxiv.org/abs/1412.6980

configuration-driven approach enables systematic experimentation and ensures complete documentation of model specifications.

Evaluation Methodology and Results

Model performance was assessed using multiple complementary metrics:

- Root Mean Squared Error (RMSE): Measures prediction accuracy in original units
- Mean Absolute Error (MAE): Provides an interpretable error magnitude
- Coefficient of Determination (R^2): Quantifies the proportion of variance explained
- Mean Absolute Percentage Error (MAPE): Expresses error as a percentage

Fig. 4.2 presents the training and validation loss curves across training epochs. This visualisation allows users to monitor the learning process and identify whether the model has successfully converged, is overfitting (divergence between training and validation loss), or requires additional training. Such diagnostic information is essential for determining model reliability and identifying potential training issues.

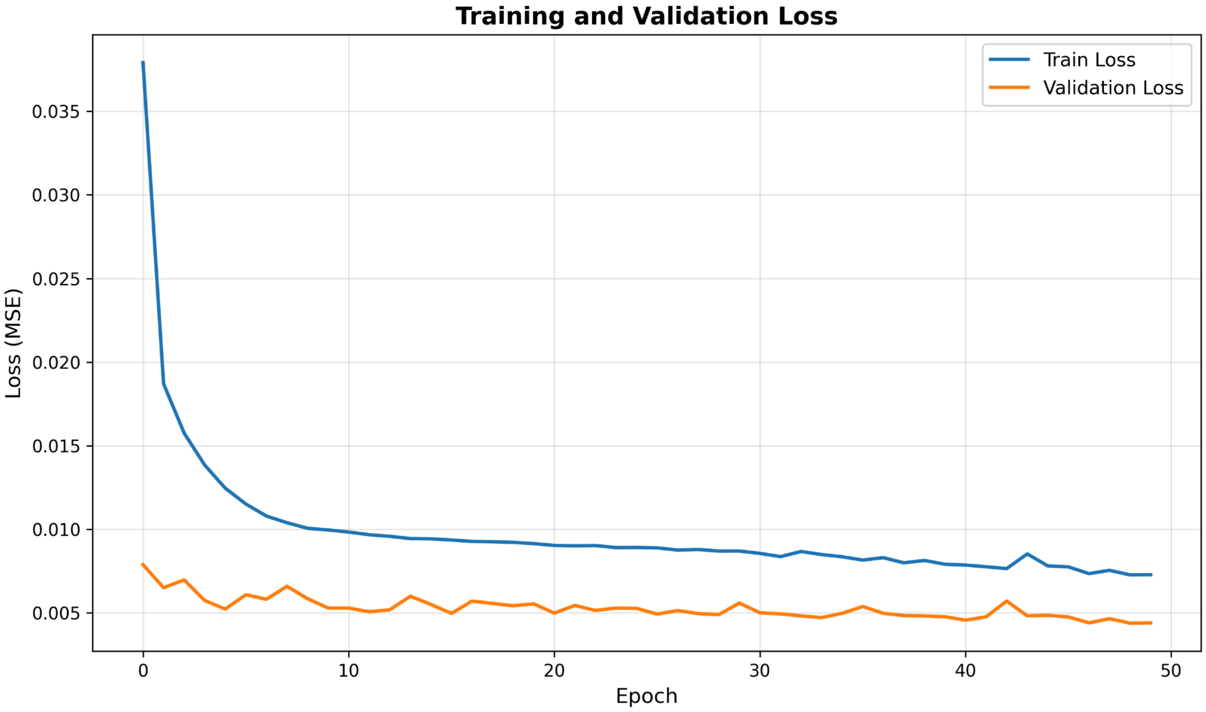

Fig. 4.2: Training and validation loss curves over epochs.

Fig. 4.3 displays model predictions overlaid with actual consumption values across the complete test set. This full-timeline visualisation enables users to assess the model's overall forecasting capability across extended periods, identify systematic biases or patterns in prediction errors, and evaluate whether the model maintains consistent performance across different time periods and consumption regimes.

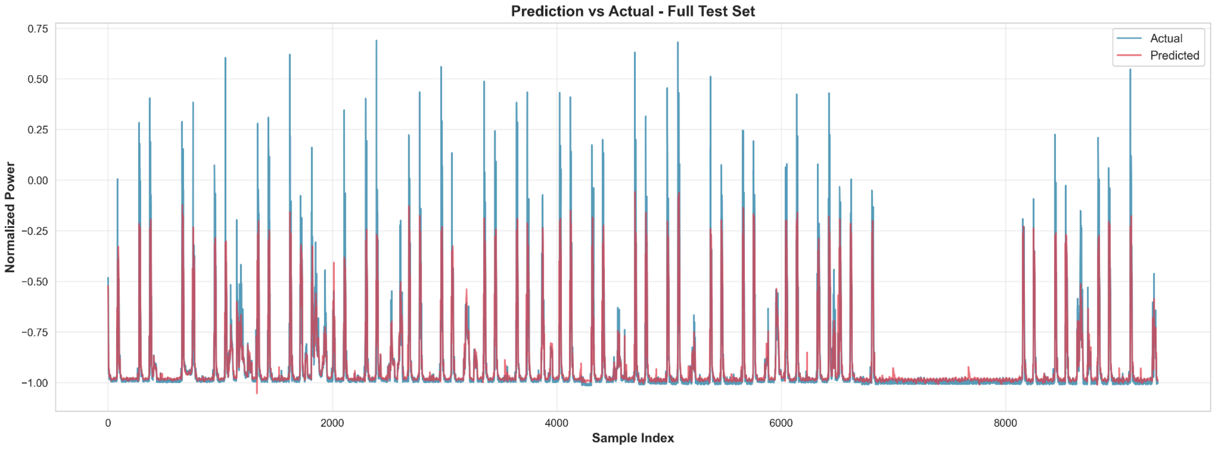

Fig. 4.3: Model predictions versus actual consumption over the complete test set.

Fig. 4.4 provides a detailed view of predictions versus actual values over a limited time window (7 days). By zooming into a shorter timeframe, this visualisation allows users to examine the model's ability to capture fine-grained temporal patterns such as daily consumption cycles, weekend versus weekday differences, and sudden consumption changes. This granular perspective complements the broader view of Fig. 4.4.

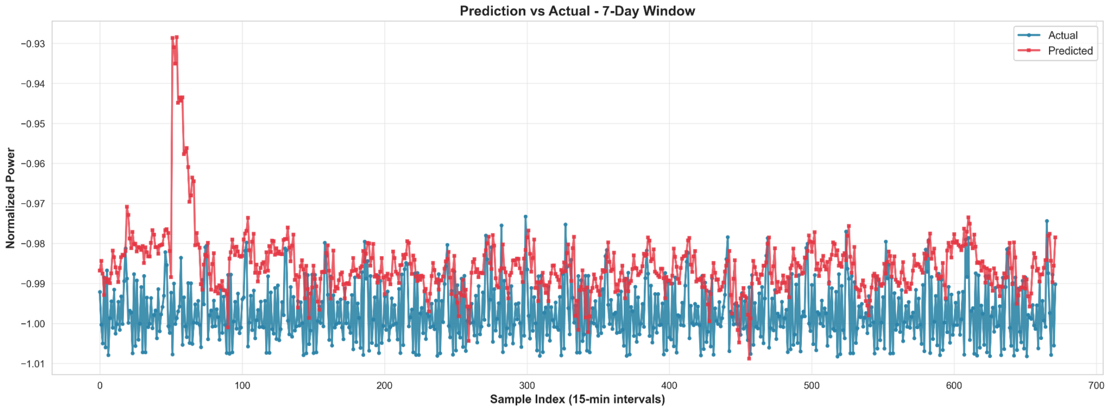

Fig. 4.4: Detailed view of a 7-day window demonstrating daily consumption cycles.

Fig. 4.5 illustrates forecast accuracy metrics (MSE and MAE) as a function of prediction horizon. This multi-step evaluation reveals how prediction quality evolves as the forecast extends further into the future. Understanding this degradation pattern is critical for end-users who need to determine appropriate forecast horizons for their specific applications and assess the model's reliability for short-term versus longer-term predictions.

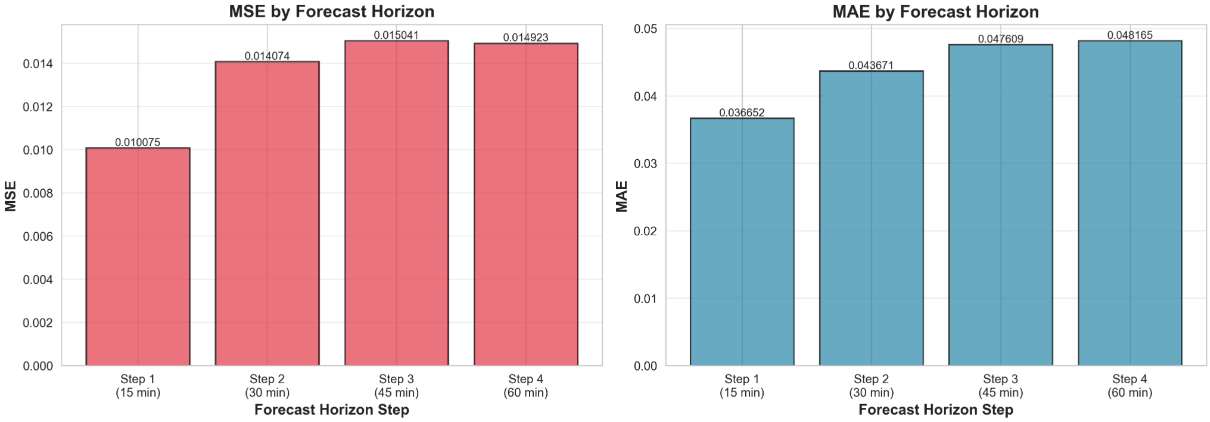

Fig. 4.5: Forecast accuracy (MSE and MAE) as a function of prediction horizon.

Fig. 4.6 presents a comparative analysis of test set performance across different experimental configurations. This visualisation enables users to systematically evaluate the impact of different pipeline configurations—such as feature selection strategies, architectural choices, or hyperparameter settings,on forecasting accuracy. By providing side-by-side performance comparisons through multiple metrics (MSE, RMSE, R²), users can make informed decisions about which configuration best suits their forecasting requirements.

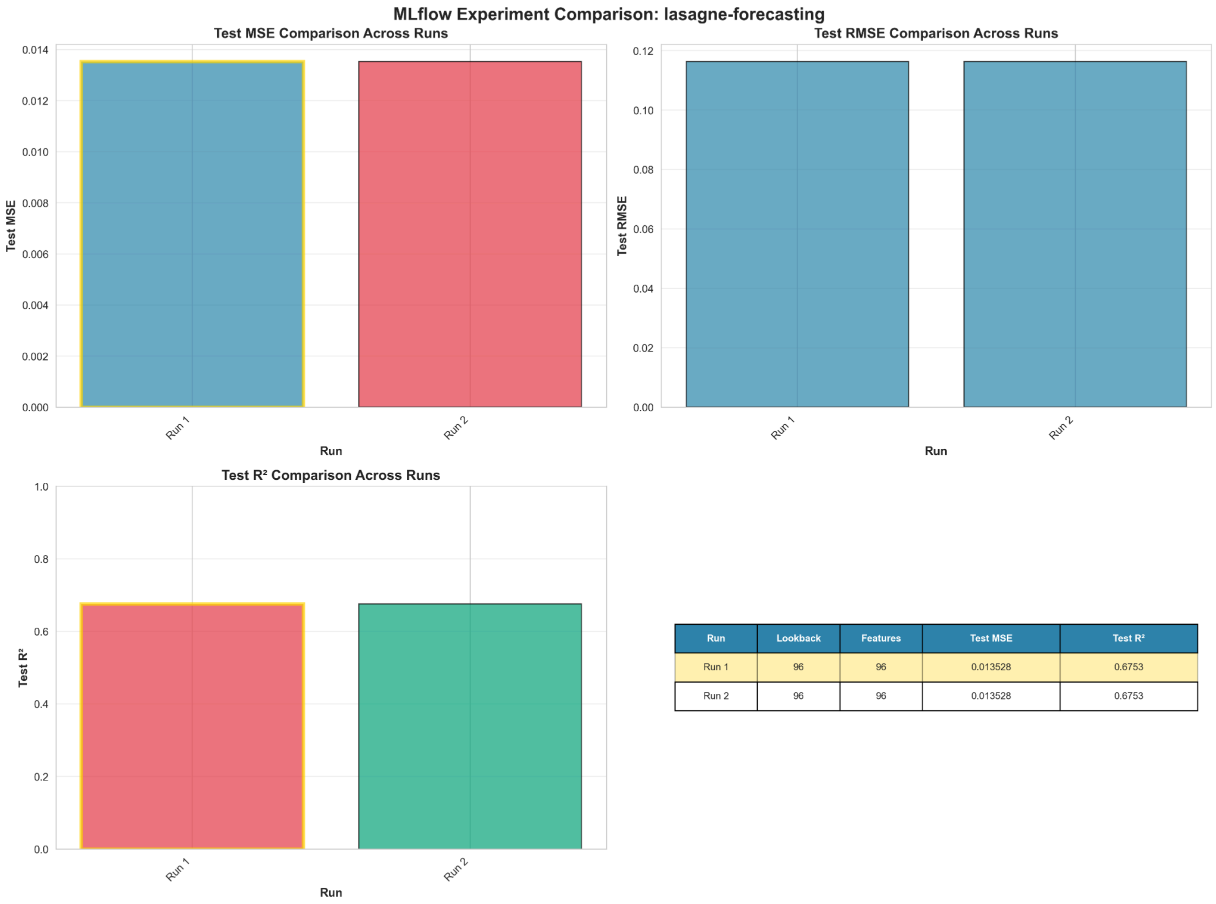

Run	Lookback	Features	Test MSE	Test R²
Run 1	96	96	0.013528	0.6753
Run 2	96	96	0.013528	0.6753

Fig. 4.6: Comparison of test set performance across experimental configurations.

4.3 Grafana visualisation Interface

Visualisations were generated to display the inference results graphically. The instances of these visualisations can be found through Fig. 4.7 to 4.9.

The Lookback (resp. Horizon) is the period before (resp. after) current time t, used in the ML training: the ML training uses a lookback period of N and a horizon of M. The weights from a ML training set are valid only for a lookback/horizon of N/M. An inference for energy consumption (prediction) at time t for Energy is made using values from t_n n = {1 : N}, resulting in M prediction curves at each time t. These are ingested to an InfluxDB and plotted as shown in the Fig. below.

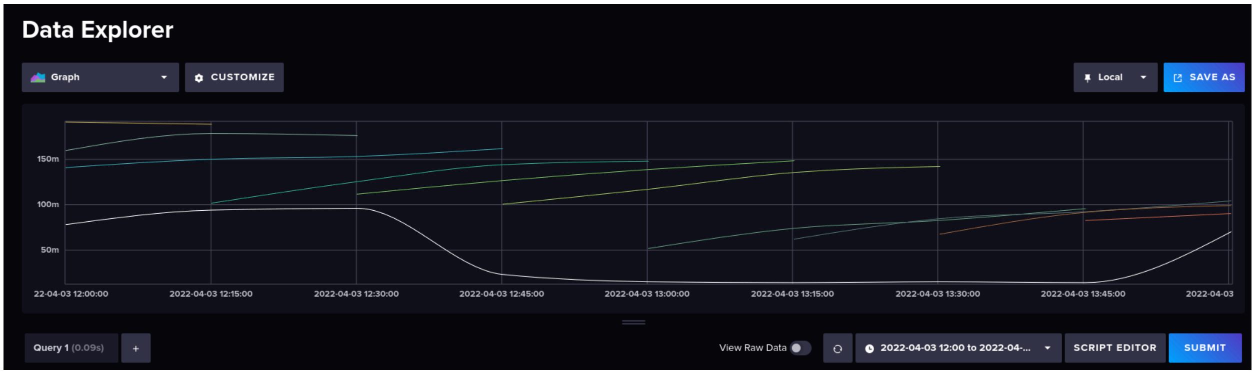

Fig 4.7: Separated prediction curves.

As can be seen below, the historical and predictions data can be plotted over time as separate plots overlaid.

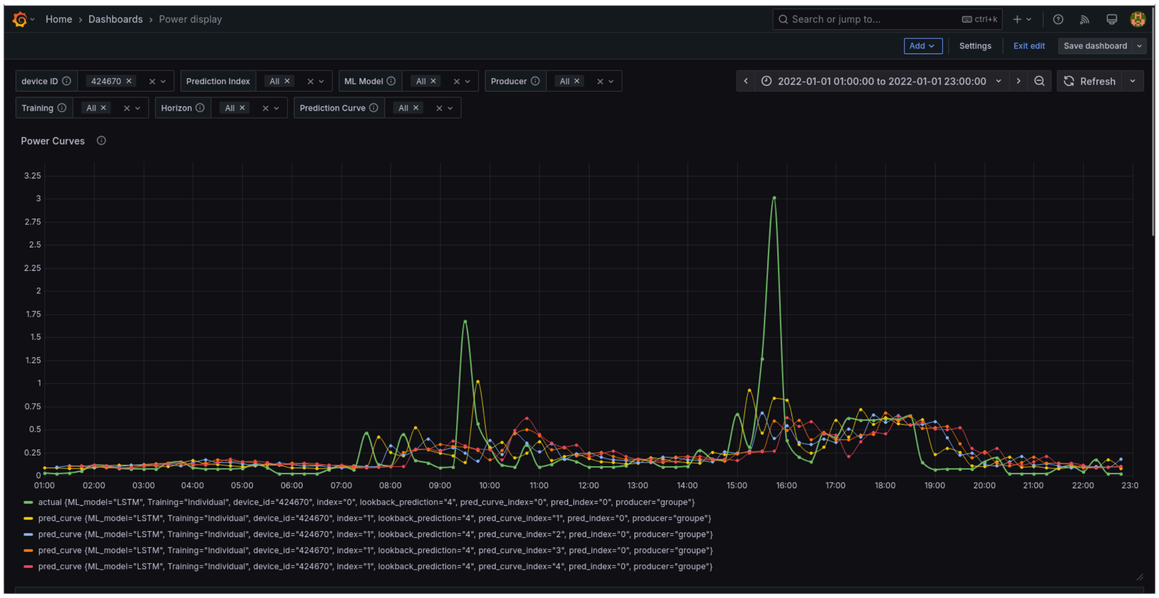

Fig 4.8: Prediction curves overlaid with historical data.

These predictions can also be averaged and overlaid on top of the historical data as shown below.

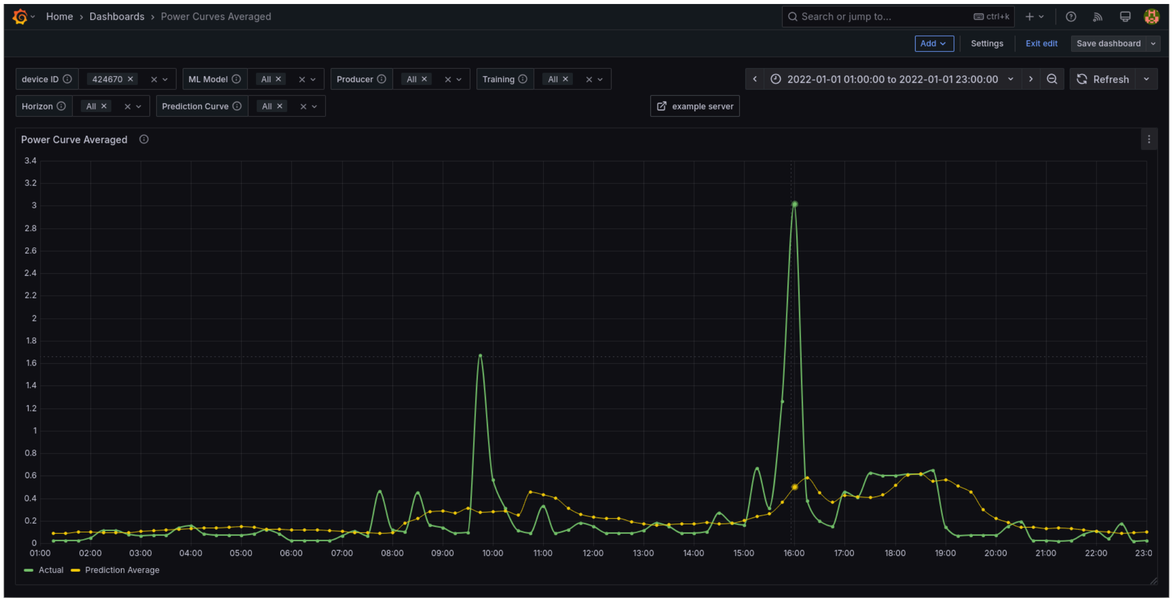

device ID: 424670

ML Model: All

Producer: All

Training: All

Horizon: All

Prediction Curve: All

2022-01-01 01:00:00 to 2022-01-01 23:00:00

Fig 4.9: Prediction curves averaged with historical data.

4.4 Conclusion

This chapter presented a configurable machine learning pipeline for electrical consumption forecasting using LSTM neural networks. Starting from the development of a data retrieval API for data acquisition from the FLOEM platform, through comprehensive data cleaning and feature engineering, to model training and evaluation. The modular, configurable architecture of the pipeline facilitates systematic experimentation and continuous improvement.

Furthermore, the visualisation through the Grafana interface has been demonstrated, enabling users to easily observe real-time data along with the predictions produced by the models.

All code, configurations, and documentation are available in the project GitLab[15] repository, ensuring that this work can be reproduced, validated, and extended.

---

[15] https://gitedu.hesge.ch/lsds

Chapter 5: The coordination model

Author: UNIGE (Giovanna Di Marzo)

5.1 Introduction

We developed and implemented a coordination model for microgrid-level energy exchange, enabling real-time, autonomous interactions between intelligent digital twins representing producers and consumers. The model is assumed to be the backbone that enables a network of GEDs (CLEMAP devices) to exchange information and improve their forecasting performance, as explained in Section 2.3 and Fig. 2.2.

It should be noted here that due to a lack of resources and time, this module has not actually been deployed on the edge CLEMAP device (Fig. 2.2). This chapter details the module's architecture. It paves the way for future deployment.

The model supports self-adaptive energy management, peak shaving, and decentralised collaborative learning through Gossip Federated and Ensemble Learning approaches. It integrates Social Acceptance criteria into digital twin algorithms to ensure user engagement and sustainability, and is validated using real and simulated data across various microgrid topologies.

We summarise here the various results and achievements linked to the coordination model and the digital twins: first, we present the coordination model; second, we describe the various digital twins acting as intelligent agents attached to the coordination and a series of algorithms aiming at regulating energy; third, we discuss a decentralised learning schema based on gossip which accommodates homogeneous models (gossip federated learning) or heterogeneous (gossip ensemble learning); forth, we present scenarios and algorithms linked to social acceptance by design, such as ensuring privacy, a dynamic pricing strategy to mitigate the tragedy of the commons, a reward schema to avoid free-riding and various common storage strategies and their impact on community fairness. These social acceptability indicators are discussed in more detail in Chapter 6, where the overall social and industrial impact is evaluated to ensure alignment between the KPIs of the pilot and the strategic objectives.

5.2 Coordination model and digital twins for smart grids

We developed a **tuple-based bio-inspired coordination model** derived from SAPERE[16], and a suite of **intelligent Digital Twins (DTs)** representing microgrid entities.

The coordination model:

---

[16] https://doi.org/10.1016/j.pmcj.2014.12.002

1. uses stigmergy to enable asynchronous, decentralised interactions among DTs via a **coordination media**, which is a shared virtual space storing and retrieving tuples of data called LSA (live semantic annotations);
2. is governed by **coordination laws** (Bonding, Decay, Spreading, Aggregation) acting dynamically on the tuples deposited in the coordination media;
3. supports **intelligent digital twins (DTs)** acting as autonomous intelligent agents - **prosumers, producers, consumers, learning twins, and regulator twins** - that exchange and regulate energy locally.

The model supports dynamic adaptation, allowing DTs to generate data and respond to real-time needs. DTs interact through the coordination platform by inserting or retrieving tuples (LSAs), thus enabling energy transactions, forecasting, and regulation. This foundational design sets the stage for exchange and regulation of energy, peak shaving and predictions (see section 6.3), for gossip decentralised learning (see section 5.4) and social acceptance by design (see section 5.5).

Fig. 5.1 depicts the various elements of the coordination model: coordination media hosting the tuples; coordination laws; and the various intelligent digital twins.

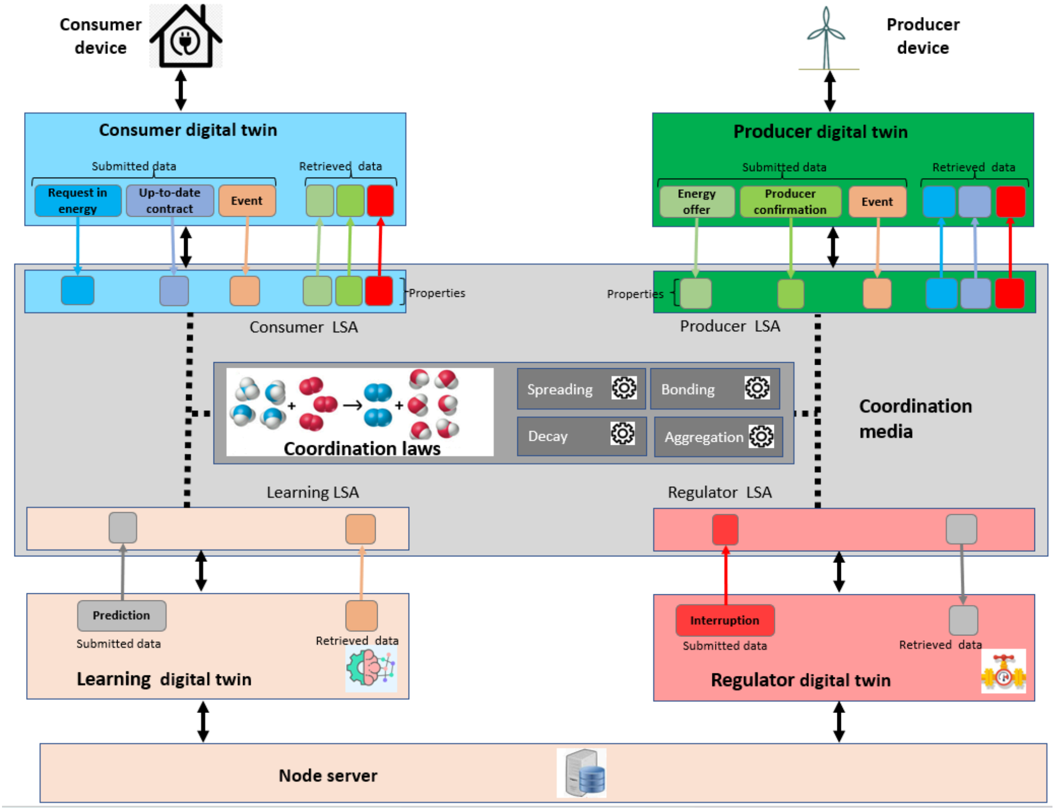

Fig. 5.1: Coordination model and intelligent digital twins

Implementation:

The **technical implementation** of the coordination platform comprises three layers: the **middleware** (managing tuple space and coordination laws), the **coordination platform service** (executing DTs and learning models), and **client applications** (web interface and simulators). DTs inherit from a generic agent class and manage energy exchanges, predictions, and regulation.

The platform supports **Markov Chains (in-platform)** and **LSTM (via external Python service)** for forecasting. REST controllers handle interactions, and **SQLite databases** store node configurations, events, and predictions. Simulators replay real and synthetic scenarios.

The system is also optimised for **low-memory devices** (e.g., Raspberry Pi), with memory leak detection, lightweight servers, and modular design. A web app visualizes node states, predictions, and learning model performance.

Deployment and data

We deployed a working coordination model on multiple edge nodes exploiting the data gathered from “Les Vergers” living lab (Meyrin deployment). The coordination model implementation supports various digital twins and algorithms that exchange information and interact together through the coordination platform. The prototype infrastructure has been set up to gather real energy usage data from *Les Vergers” and run the experiments. The optimised low-memory version for Raspberry Pi has been provided to be used in conjunction with the other LASAGNE components.

5.3 Exchange and regulation of energy among digital twins

As said above, we developed several intelligent digital twins (DTs) working on behalf of consumers, producers, learning and regulator devices (peak shaving controllers):

- **Consumer and producer DT** interact to perform contracts regarding the amount of energy exchanged (no pricing support yet), so as to optimise energy use.
- **Learning twins** forecast production and consumption locally, through various techniques (Markov chain, LSTM).
- A **regulator twin** controls the energy levels and stops any producer or consumer twin in case of peaks (anticipating or reacting on actual data).

We designed and implemented several **algorithms and mechanisms** enabling DTs to coordinate their actions in order to autonomously manage energy exchange and regulation.

Contract generation algorithm

The system provides a **contract generation algorithm** where consumer DTs request energy and producer DTs respond with offers, leading to dynamic contract formation and validation. The system supports urgent requests, adaptive responses to fluctuating demand. The steps to establish contracts are as follows (Fig. 5.2). Step 1: sending of the request by the consumer: The consumer twin submits its request through its LSA to report its need for energy supply. Step 2: sending of offers by producers: The producers retrieve the new consumer requests from their LSA. Once the requests have been retrieved, each producer then generates new elementary offers to meet the different requests. A producer applies a prioritisation policy to choose the requests to handle first. For example, a policy can classify the requests by decreasing the level of urgency, then by increasing power. The producer generates offers within the limit of its available wattage. As long as its remaining wattage is positive, it generates a new offer, even if it cannot cover all the requested wattage. Step 3: issuance of a new contract by the consumer: The consumer retrieves the offers that have been issued by the producers. Periodically, it checks whether all its retrieved offers meet the need. Similarly, it applies a policy that determines the offers to choose as a priority. The consumer then goes through the different offers in an iterative way, following the chosen policy: as soon as the offer aggregated (from the elementary offers) meets the need, the consumer issues a new contract based on the selected offers. Step 4: confirmation of a contract by the producers concerned: To have the contract confirmed by the producers, the consumer twin submits it in its LSA. Then, each producer involved retrieves the contract and checks whether it agrees or not with the terms and then submits its confirmation or refusal in its LSA: in this way, the consumer will retrieve the producer confirmation or refusal and will update the status of the contract validation. Step 5: validation of a contract by the consumer twin: Each time the consumer twin retrieves one confirmation, it updates the contract status to include the producer's approval (or disapproval), and then, submits the new contract content to its LSA so that all stakeholders involved have the updated version. The consumer validates the contract only when all the producers involved have responded positively. If so, all stakeholders involved retrieve a validated contract: the contract then takes effect, and the producers can start supplying energy. If any of the stakeholders disapprove the contract, the consumer cancels it, and a new bidding cycle begins immediately.

Predictions

The system provides **predictions** of production and consumption. We implemented two variants, through Markov Chain and through the LSTM model. The learning twin uses Markov Chains model to predict the trend of each power variable at the node level. This learning model is a stochastic model, which uses transition matrices to define the transition probabilities of all possible states. This model has the advantage of being able to predict time series fluctuating over time, thanks to the use of a sliding window: the latter allows the model to be re-trained by focusing on the data of the last N days. Furthermore, this learning model is relatively easy and quick to implement, which facilitated its integration into the coordination platform. We use the LSTM learning model as an alternative to Markov Chains to predict the power variables of the node state. Indeed, LSTM is known to provide excellent performance for time series learning and to be less sensitive to high variations in time intervals. Given that these two models are implemented, the learning twin can use either Markov Chains or LSTMs, depending on the settings chosen in the configuration of the coordination platform.

Peak shaving

The system also provides **peak shaving** through proactive regulation. Peak shaving aims at providing stable conditions for the grid. Decisions are taken either based on actual grid conditions or based on predictions. The regulation is carried out entirely by the regulator twin (and not by an external or manual system). The regulator twin acts locally and its impact is limited to the other Digital Twins (consumer or producer twins) connected to the same node. At regular intervals, the regulator twin retrieves the node global state that the learning twin recently stored in the database. From the different totals, the regulator twin detects the production and consumption flows that exceed the alert threshold. If the total produced or the total consumed electricity exceeds the threshold, the regulator twin applies a policy to determine the DT to stop first so that the node can get back under the alert threshold. The regulator twin thus obtains an ordered list of DT to stop. For example, the regulator twin can stop by order of priority the producers with the highest power. The regulator twin stops one by one the Digital Twins contained in this list, as long as the updated total remains above the alert threshold. The learning twin predicts node states using Markov Chains and LSTM models, while the regulator twin uses these forecasts to prevent overproduction or overconsumption.

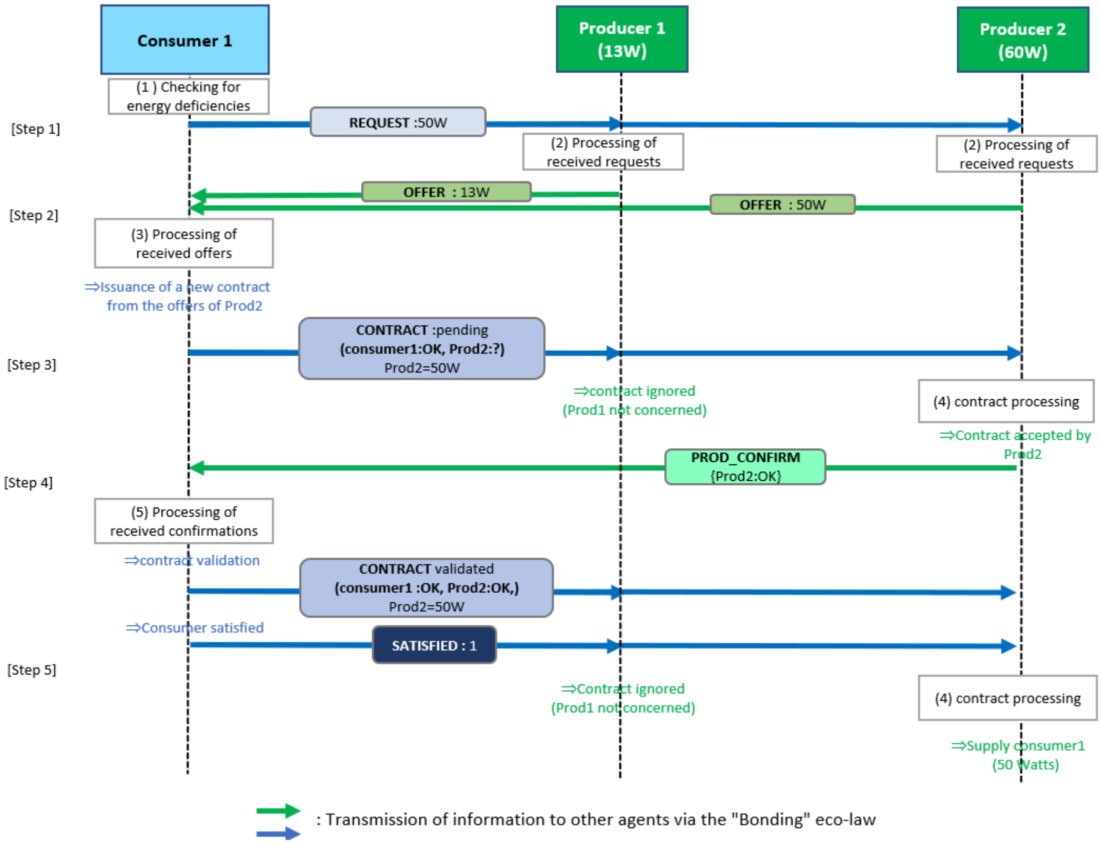

Fig. 5.2. Contract generation in the form of a sequence diagram. The different steps described above appear on the left.

Results

Evaluation metrics include **response time** and **failure rate**, tested across various scenarios (technical, realistic, degraded, and living-lab). Results show that prioritization policies outperform random ones, and the coordination model effectively supports self-adaptive energy management.

5.4 Gossip coordination mechanism for decentralised learning

We developed a **Gossip-based coordination mechanism** for decentralised machine learning across microgrid nodes. Decentralised knowledge sharing between different nodes in a network can be achieved in different ways, depending on whether the learning models have the same structure or not. Either the different entities manage learning models with the same structure, in which case it is possible to exchange model weight parameters to make each node's model ”more efficient” in predicting electrical behaviour at a local level or cluster level. In this case, we can try to adopt a Federated Learning strategy based on the gossip mechanism. Either this is not the case: the entities have heterogeneous structural models, and exchanging model parameters makes no sense. In this case, it is possible to benefit from the learning experience of other nodes by exchanging the prediction data generated by the different nodes. It is always possible to aggregate prediction results, whatever the learning model used by the nodes. In this situation, we can adopt an ”ensemble” learning strategy, based on a gossip mechanism. In both cases, we are talking about decentralised learning using the gossip mechanism, but in the first case, we are talking about Federated Learning and applying aggregation to learning models, while in the second case, we are talking about Ensemble Learning and applying aggregation to the predictions themselves. We implemented both approaches.

Gossip Federated Learning

Nodes exchange and aggregate homogeneous learning models (e.g., LSTM, Markov Chains). As shown on Fig. 5.3, periodically, each node: Step 1. Trains iteratively its own learning model using its private local data updated with the latest observations (see step number 1 on the upper figure). Step 2. Receives model weights from surrounding nodes which are directly linked (see step number 2 on the bottom figure). Step 3. Aggregate the models received into one merged model. Its own model is also included in the aggregation (see step number 3 on the bottom figure). Step 4. Spreads its updated model to the surrounding nodes which are directly linked (see step number 4 on the bottom figure).

Gossip Ensemble Learning

Nodes exchange prediction results from heterogeneous models. The general mechanism remains the same as the one for Gossip Federated Learning. Indeed, the general stages (1, 2, 3, 4) which involve local updates, reception, aggregation, and spreading are similar to those of Gossip Federated Learning. The difference is that these operations are applied to predictions rather than learning models.

Results

The gossip mechanism combines **aggregation and spreading**, enabling scalable, privacy-preserving learning. Aggregation operators include **sampling size**, **power loss**, **min loss**, and **profile similarity**. Experiments on the Les Vergers dataset show that Gossip Federated Learning with LSTM outperforms local learning. Gossip Ensemble Learning performs well at the cluster level but less so at the node level due to model heterogeneity. Bandwidth optimisation techniques (e.g., compression, synchronization) enhance efficiency.

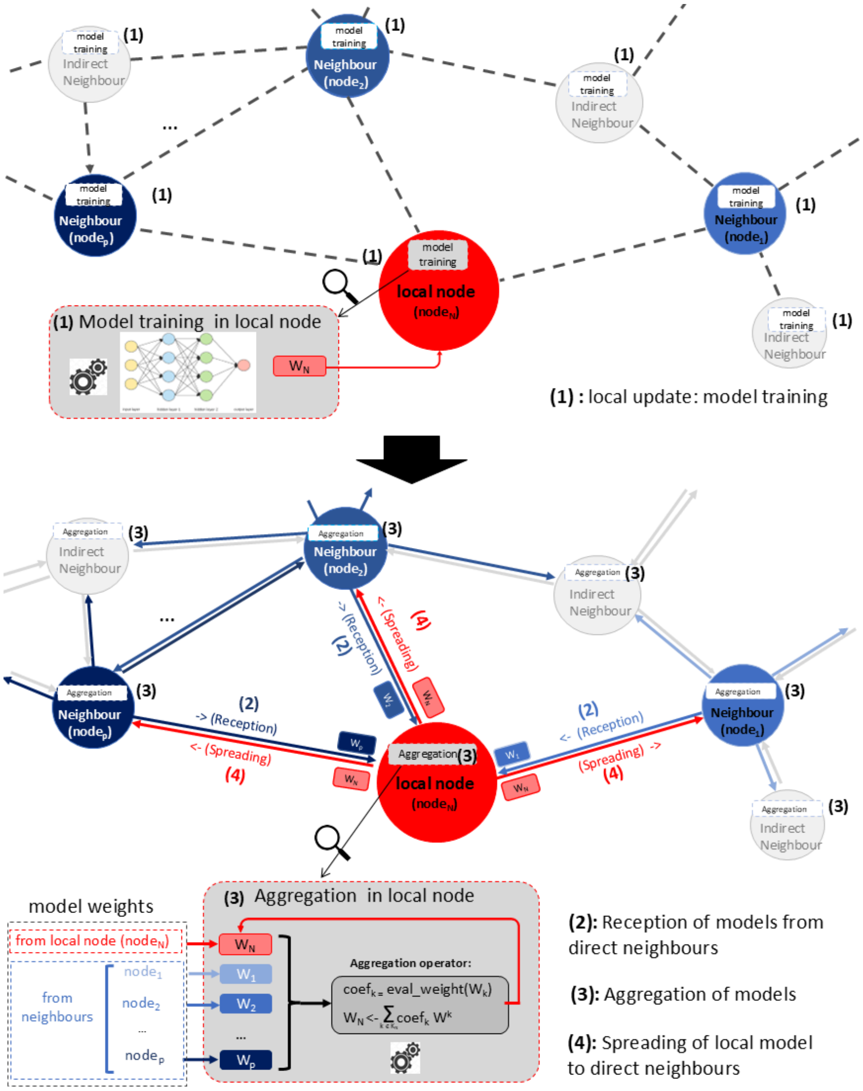

Fig. 5.3. Using coordination mechanisms to implement the Gossip Federated Learning Approach (Steps 1 to 4).

5.5 Social acceptance by design

In collaboration with HES-SO Valais we integrated **social acceptance criteria by design** into the behavior of Digital Twins, addressing user concerns in microgrid adoption. Based on studies by HES-SO Valais, key issues include **privacy**, **free-riding**, **tragedy of the commons**, and **community fairness**. DTs are adapted to: (1) Protect personal data via wrappers; (2) Implement **dynamic pricing** to discourage overuse; (3) Assign **reward credits** based on prosumer contributions; (4) Manage **energy storage** (private vs. common) to ensure equitable access.

Privacy

To meet social acceptance criteria, we apply confidentiality rules to the data exchanged in the tuple space: personal information about a Digital Twin cannot be visible to its peers. For example, the different supplies made by a producer number A are not visible from another producer number B. These confidentiality rules apply to the various data exchanged such as requests, offers, contracts or producer confirmations: they consist in protecting each data submitted through a LSA individually. Access to ”personal” properties is only permitted for the concerned DT. We use the notion of ”wrapper” to protect access to sensitive objects. We therefore declare the sensitive object as a private attribute in a wrapper class which contains it, and which manages access. The access method checks whether the Digital Twin has correctly authenticated to the microgrid and, if so, it determines the extent of visibility the DT has on the requested data. The method restricts the data to return, according to the obtained visibility.

Tragedy of the commons and dynamic pricing

The **tragedy of the commons** scenario corresponds to the case where the best individual decision for a given Digital Twin turns out to be the worst-case scenario for them as a group. This is a classic scenario which causes a shortage of electricity on a microgrid scale. We therefore propose to study this scenario with the aim of highlighting such a risk and providing a solution through an alternative scenario. As an example, let’s consider the charging of electric vehicles (EV) at night when energy prices are less expensive, for instance between 11 pm and 5 am. Individually, it is a good policy to start charging the EV from 11 pm onward. However, if everybody in a neighbourhood takes this same decision, this leads to a peak of energy consumption at 11 pm.

To address this issue, we designed and implemented a **dynamic pricing policy**, where the price depends on the real-time demand. To do so, we have defined a price escalation factor that is a function of the remaining availability of the producer in watts. This factor makes pricing dynamic, given the current increase in demand: if the producer’s availability decreases, the price increases in real time, depending on the producer’s real time availability. The resulting pricing therefore encourages unserved consumers to shift their demand for energy to another time slot when the producer is more available. In this way, we can spread the demands of different consumers over time using a dynamic pricing which favours the least popular periods. A new iteration is carried out when there are unfulfilled requests at the current time. At each iteration, the DTs perform the following operations: 1. Each producer supplies requests according to its capacity. This reduces its available capacity to zero if the overall capacity is not sufficient to meet all requests. 2. Each producer applies dynamic pricing by recalculating all slot tariffs based on updated availability. 3. Each dissatisfied consumer updates the tariff table and revises the time slot of its request according to the new tariffs (it must target the period corresponding to the lowest price).

Free-riding and awards computation

For a microgrid community to function, the various prosumers are supposed to share energy resources and therefore contribute to supplying their neighbours when they are short of energy, in the fairest possible way. Indeed, a failure to respect this principle of equality of supply leads to frustration and demotivation, and therefore to a reduction in the general well-being of the community. We call **free riding** behaviour this type of behaviour consisting of taking advantage of the services offered by the community while ceasing to contribute. To tackle these failures of fairness principles, the adopted strategy consists in discriminating or penalising prosumers through penalties in the decision-making process. This strategy can be integrated into the algorithms executed by the DTs. They can interact in a way to penalise any free riding type behaviour by detecting it on the surrounding prosumers; or by adapting their own supply strategy if this behaviour is detected on a peer. We propose two variants that can be adopted by a prosumer in response to the free riding behaviour of another prosumer: (1) no longer sending an offer to the free riding prosumer; (2) applying a penalty in the energy price per kWh for the offers sent to free riding prosumers.

For data privacy reasons, we defined that the regulatory agent evaluates the behaviour of the prosumer twins attached to the same node and assigns them reward credits (or penalties) according to the obtained evaluation. Each prosumer twin will then emit its current reward credit in the energy requests, which will allow the producers to adapt their supply strategy according to the reward credit. Only the final score shown in the energy requests will be visible to the other peers. The regulator twin evaluates and provides reward credit to each prosumer according to the 3 following criteria of evaluation (Fig. 5.4): (1) Level of altruism: distribution of energy surplus to prosumers in the community when the surplus exceeds a certain threshold; (2) Ability to save energy: avoid excessive consumption in relation to current availability at the node level; (3) Equity in distribution (for the same level of award): disparity in requests responding by ”client” prosumers ratio of served request response time; (4) Ratio of ”green energy” consumption. We have aligned our algorithms on sociological studies made in the framework of the LASAGNE project and on existing works to handle free riding behaviours. For example, if a prosumer Digital Twin has an individualistic behaviour, such as profiting from other Digital Twins without contributing to the community, it will be rated with a low score. As a result, each peer prosumer twin will penalise the ”profiting” twin according to its policy applied: the “free rider” will be either excluded as a requester of energy, or it will face an important rise in price per kWh, or it will become less of a priority as a potential consumer.

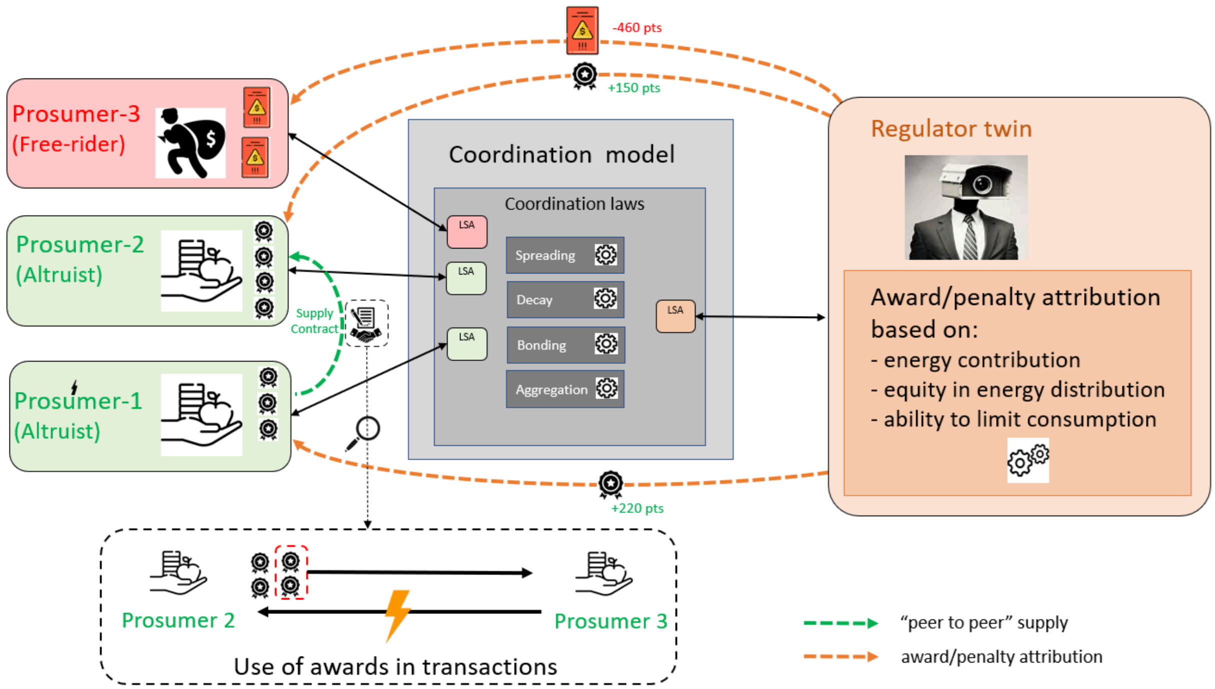

Fig. 5.4: Reward scenario with 2 altruistic prosumers and 1 'free rider' prosumer

Community fairness and storage management

In a microgrid, the prevention of energy shortages is one of the main concerns and responds to a large number of sources of worry, particularly when certain supplies are critical. To meet this challenge, we propose to model energy storage at the level of a microgrid and to integrate storage management into the behaviour of digital twins. We consider local energy storage as a way of coping with a general shortage that may be caused by the temporary shutdown of an external supplier. More concretely, we propose two different solutions in terms of organisation and implementation: (1) either each prosumer manages its own storage exclusively for its own needs, or (2) all the prosumers in the microgrid share a common storage that can be used by all the prosumer members of the microgrid.

Results

Tragedy of the commons: Various experiments involving various shortage cases show that dynamic pricing efficiently tackles the tragedy of the common.

Free riding: Results confirm that a prosumer twin can exclude as a potential consumer any prosumer penalised for bad behaviour. This could encourage a "free riding" prosumer to cooperate further by supplying energy to the other prosumers requesting energy. Experiments with the rewards strategy confirm that the use of reward credits has a considerable impact on energy prices. This could be an effective way of penalising "free riding" behaviour and encouraging prosumers to improve their involvement in the smooth running of the microgrid community.

<u>Community fairness</u>: We ran several scenarios involving energy stored locally or commonly. Results show that common storage significantly reduces shortages and improves fairness.

5.6 Conclusion

In summary, we developed various elements as part of the LASAGNE framework. (1) A Coordination Model for Smart Grids - a bio-inspired, tuple-based coordination model was designed and implemented, enabling Grid Edge Devices (GEDs) to interact and collectively manage energy exchanges and regulation. (2) Intelligent Digital Twins (DTs) to represent microgrid entities (producers, consumers, regulators, learners). These autonomous agents coordinate energy transactions and adapt to real-time grid conditions. The system integrates Multi-Agent Systems and self-organising principles for dynamic adaptation. The implementation of the coordination model and intelligent DTs includes as well a lightweight implementation providing a memory-efficient version of the coordination platform to run on low-resource devices like Raspberry Pi, enabling real-world deployment along the other LASAGNE components; (3) Decentralised Learning via Gossip Mechanism - A gossip-based coordination mechanism was implemented to support decentralized learning across nodes. This includes both Gossip Federated Learning and Gossip Ensemble Learning, improving prediction accuracy and scalability; (4) Social Acceptance by design - The system incorporates social criteria such as privacy protection, dynamic pricing, and reward mechanisms to discourage free-riding and promote fairness.

As limitations and future perspectives we can mention the following points. (1) Integrate Price Negotiation: Enable peer-to-peer and market-based energy trading with predictive pricing models. (2) Improve Physical Integration: Enable DTs to send control signals to smart meters (actuators), completing the feedback loop. (3) Scale Deployment: Move from small-scale simulations to real-world infrastructures with dozens of distributed nodes. (4) Incorporate Meteorological Data: Use weather variables to improve energy forecasting models. (5) Further integration of social acceptance: extend DTs to incorporate **salient attributes** (e.g., grid connection, energy mix, data sharing) into decision-making.

This work demonstrates that distributed, adaptive systems using coordination models and intelligent DTs are well-suited for managing the complexity and social acceptance of smart grids. Unlike centralised systems, this approach supports local autonomy, scalability, and resilience, paving the way for sustainable and socially accepted energy communities.

References and research outputs

PhD:

- GLASS, Philippe. *Coordination model and Digital Twins for managing energy consumption and production in a smart grid*. Doctoral Thesis, 2025. doi: https://doi.org/10.13097/archive-ouverte/unige:186221

Journals:

- Glass, Philippe and Di Marzo Serugendo, Giovanna: **Coordination Model and Digital Twins for Managing Energy Consumption and Production in a Smart Grid.** November 2023. https://www.mdpi.com/1996-1073/16/22/7629
- Glass, P., & Di Marzo Serugendo, G. (2025). **Gossip Coordination Mechanism for Decentralised Learning**. *Energies*, *18*(8), 2116. https://doi.org/10.3390/en18082116

Chapter in book

- Glass, P., & Di Marzo Serugendo, G. (2026). **Intelligent digital twins for social acceptance by design in smart grids.** Chapter 7 In Digital Twin Technology and Smart Grid. Hassani, H., Bagula, B.A., Varlamis, I. (Eds). Elsevier. (to appear)

Posters

- Di Marzo Serugendo, Giovanna and Glass, Philippe: “Lasagne - implementation smartgrid in the form of a coordination system of digital twins “, October 2022- poster presented at the data-science day, September 15, 2022. https://archive-ouverte.unige.ch/unige:164521
- Di Marzo Serugendo, Giovanna and Glass, Philippe: “Coordination model and digital twins for managing energy consumption and production in a smart grid“, October 2023. - poster presented at the PdD students day, September 2023 https://archive-ouverte.unige.ch/unige:172075

Open source:

- Sources of the implementation available in free access, in a GitHub Repository (https://github.com/philippe-glass/energy/tree/master/).
- A docker image of the coordination platform is available on docker hub. (https://hub.docker.com/repository/docker/philippeglass1/coordination_platform/general)

Chapter 6: Social acceptability

Author: HES-SO//Valais (Emmanuel Fragnière)

6.1 Introduction

A successful transition to smart micro-grids is essential for integrating renewable energy and for consumers to become active producers (or 'prosumers'). This requires more than just technological prowess. It requires social acceptance by design. Without the full commitment and engagement of users in the smart grid community, the project’s chances of success are minimal. This chapter outlines the structured approach developed within the LASAGNE project, which defines a comprehensive toolbox used to identify, quantify and optimise social acceptance factors in the design of micro-smart grid energy communities. This evaluation is based on the social acceptability indicators defined in Chapter 5 (e.g. prosumer profiles, user comfort improvement and community engagement). These elements ensure traceability between social perception measures and industrial impacts, in line with the KPIs set out in the Grant Agreement.

The toolbox moves progressively from qualitative investigation and knowledge structuring to quantitative measurement and, ultimately, optimisation modelling. This ensures that the final system configuration maximises user engagement and cooperation.

6.2 Structure of the Social Context

The first stage of developing the toolbox involved understanding the social context and latent needs of potential microgrid users and stakeholders. This foundation was established through ethnographic research, primarily comprising qualitative fieldwork, such as semi-structured interviews and focus groups, with Swiss smart grid stakeholders. These stakeholders included residents of eco-districts and experts.

The aim of this qualitative research was to gain a deep, nuanced understanding of the social and cultural factors that influence adoption and acceptance, rather than to achieve statistical representativeness. The research revealed the key attributes influencing acceptance across various dimensions, such as energy infrastructure, pricing mechanisms, information sharing, community engagement, regulations and governance.

Translating the key findings from this exploration into potential drivers and barriers revealed that restraining forces were often more influential than driving forces (see Fig. 6.1).

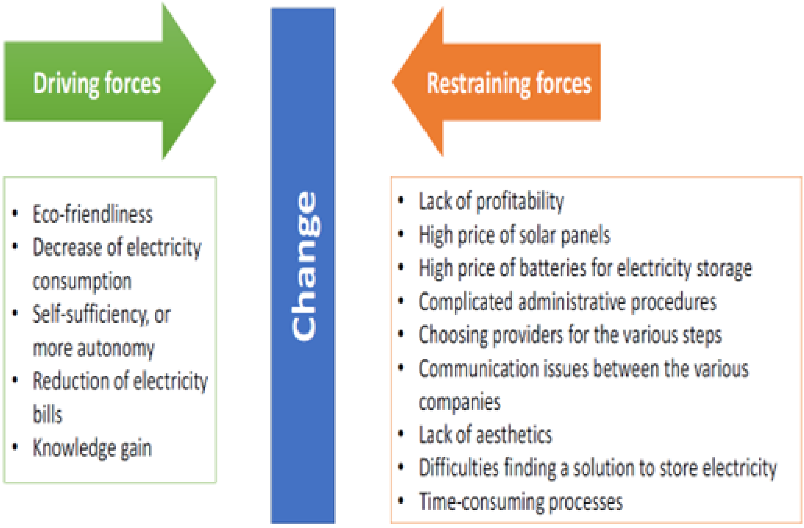

Fig. 6.1: Force Analysis Diagram – Social Acceptability of Microgrids

6.3 Structuring Social Insights: Preprocessing Tacit Knowledge

In order to transform the rich, complex insights gathered during the ethnographic phase into actionable design elements, it was necessary to pre-process this qualitative, or tacit, knowledge in order to structure it. This structuring was accomplished using the SECI (Socialisation, Externalisation, Combination and Internalisation) model of knowledge creation.

The SECI model was used to interpret how a community absorbs and shares knowledge regarding energy management and sustainability practices.

- **Socialisation** relies on tacit knowledge shared through direct interactions (e.g. workshops), fostering trust and shared values. Key findings related to this phase included the necessity of group cohesion and shared values, such as eco-friendliness, as well as a strong sense of responsibility and interdependence among prosumers.

- **Externalisation** occurs when tacit knowledge is converted into explicit knowledge; for example, when users interact with smart meters and receive real-time data about their consumption behaviours.

- **Combination** involves AI algorithms processing this explicit data (e.g. optimising energy storage or demand response).

- **Internalisation** is when prosumers absorb the explicit knowledge provided by the smart grid and convert it back into adapted behaviours and habits.

This process transforms abstract social concerns (e.g. the need for 'trust') into concrete, measurable attributes (e.g. the need for 'regular monitoring' or 'data access') that can be used in the subsequent design phase.**Designing Prosumer Scenarios: Creating Profiles for Measurement**

The salient, structured attributes resulting from the qualitative and SECI analyses were then used to develop card-based surveys for quantitative measurement. This technique is central to conjoint analysis, enabling respondents to intuitively evaluate trade-offs between different system attributes.

Based on iterative focus group feedback, the attributes were simplified and organised to avoid cognitive overload and ensure clarity for non-specialists. The qualitative research initially produced a binary scale (option A or B, see Fig. 6.2), though the final card-based survey used six profiles (cards) for ranking preference.

The key attributes representing the different dimensions of microgrid system design and their corresponding levels were defined.

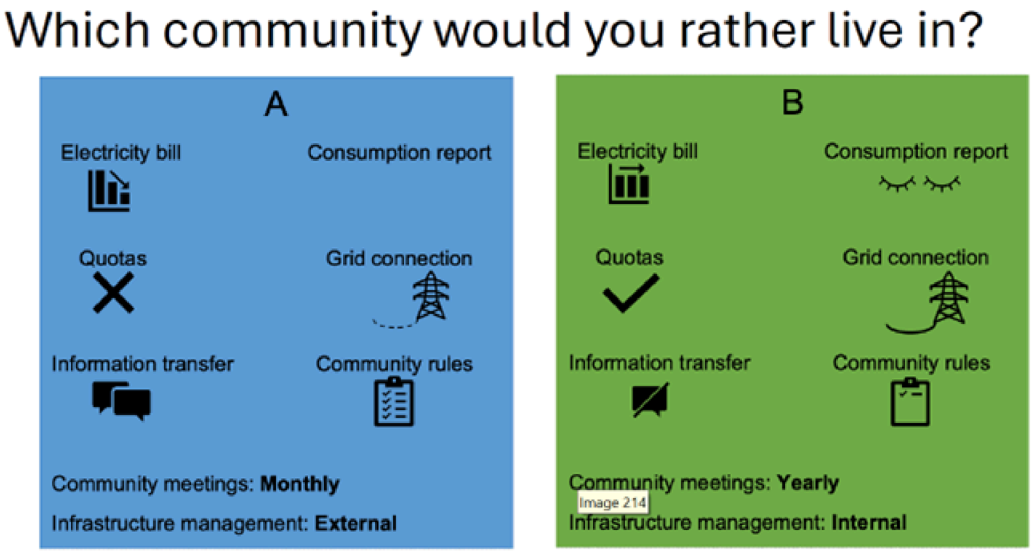

Fig. 6.2: Examples of first prosumer profiles that were created based on the ethnographic survey

6.4 Quantifying Preferences: Feature Valuation via Conjoint Analysis

Once the cards representing the scenarios had been defined, Conjoint Analysis (CA) was employed to quantify the relative importance of these social and technical features. Respondents were asked to rank the card profiles in order of preference, from most to least preferred. This has been done with a final prototype of 6 cards and programmed in the Sphinx market analysis software.

Conjoint analysis quantifies these preferences, yielding part-worth utilities (i.e. econometric models) for each attribute level. These utilities reveal how consumers prioritise different features when making decisions about microgrid participation. For example, they might balance the preference for autonomy (photovoltaic) against complexity (dynamic pricing) or community participation (frequent meetings).

6.5 More about social salient attributes

Behavioural features (governance and community):

These features relate directly to the social rules and governance within the energy community. The ethnographic phase emphasised the importance of community involvement, as well as raising concerns about over-regulation.

- Quota: Yes/No

Meetings: Frequent/Rare (addressing community engagement and governance complexity).

Technical configuration features (infrastructure):

These features address the physical setup and autonomy of the system, both of which are crucial factors influencing trust and perceived resilience.

Electricity: External/photovoltaic/both (covering grid connection preferences).

- Battery: With/without (related to storage and self-sufficiency; often mandatory if relying only on photovoltaic).

Economic features (pricing model):

The pricing mechanism is a fundamental attribute that directly impacts the monetary costs and perceived benefits for the prosumer.

- Pricing: Flat/Dynamic (Dynamic pricing adapts to supply and demand, requiring a higher level of commitment and potentially achieving a higher level of social acceptance).

6.6 Optimal Design of Prosumer's configuration for Energy Community

In order to accurately measure social acceptance, the preference for the new smart grid design must be compared to a standard alternative scenario. In this study, the alternative choice was defined as connection to the external electrical grid without specific efforts towards reducing emissions. This alternative design is represented by a specific utility value and serves as the reference point against which social acceptance is calculated (the weighted proportion of users who prefer the new design).

Tracking results: Part-Worth Utility Quantification

The quantified preferences (part-worth utilities) generated by the conjoint analysis serve as critical inputs for the subsequent optimisation phase. These utilities are necessary to determine whether a specific respondent prefers the proposed new microgrid design to the alternative design.

The Optimisation Engine: The Share-of-Choice Model

The final step in the social acceptance toolbox is the share-of-choice optimisation model. This meta-model combines measured social preferences (utilities) with a synthetic production model characterised by cost structures and sustainability goals (CO_2 reduction).

The objective of this optimisation model is to determine the most cost-effective microgrid configuration that maximises social acceptance while minimising costs and meeting sustainability targets. This approach recognises that technical constraints and social acceptability must be balanced for the successful long-term diffusion of microgrids.

The model formulation is structured as a 0–1 programming problem, in which decisions regarding the microgrid configuration are made to minimise costs, subject to constraints regarding social acceptance and minimum CO_2 reduction.

Key parameters of optimisation (the design variables):

The variables driving the optimisation are divided into two categories:

1. Service configuration variables (X(i, j)): These represent design decisions. X(i,j) = 1 if attribute i is set to level j (e.g. battery is set to 'without'), and 0 otherwise. Constraints ensure that only one level is chosen per attribute.
2. Preference variables (Y(k)): These measure social acceptance. Y(k) = 1 if respondent k prefers the new design to the alternative and 0 otherwise. Total social acceptance (A) is the weighted sum of these preferences.

The constraints also enforce physical and social rules mathematically, such as the necessity of having a battery and consumption quota if the community is solely photovoltaic and not connected to the external grid.

6.7 Conclusion

This methodology demonstrates how to construct a progressive toolbox for achieving 'social acceptance by design' within smart micro-grids, moving systematically from qualitative community research to quantitative design optimisation.

Coupling ethnographic exploration with the SECI model and translating insights into card-based conjoint analysis surveys allows us to feed the resulting part-worth utilities into a Share-of-Choice optimisation model. This provides a robust framework for identifying system configurations that balance environmental impact, economic feasibility and user commitment.

For example, numerical experiments using fictitious data revealed an optimal design with a high acceptance rate (70%), competitive annual costs (CHF 19,800) and commendable CO_2 reduction (7.7 tons/year).

The specific characteristics of this optimal, balanced design were:

- Electricity source: Both (grid + photovoltaic)
- Battery: Without
- Pricing: Flat
- Quota: No
- Meeting frequency: Rare

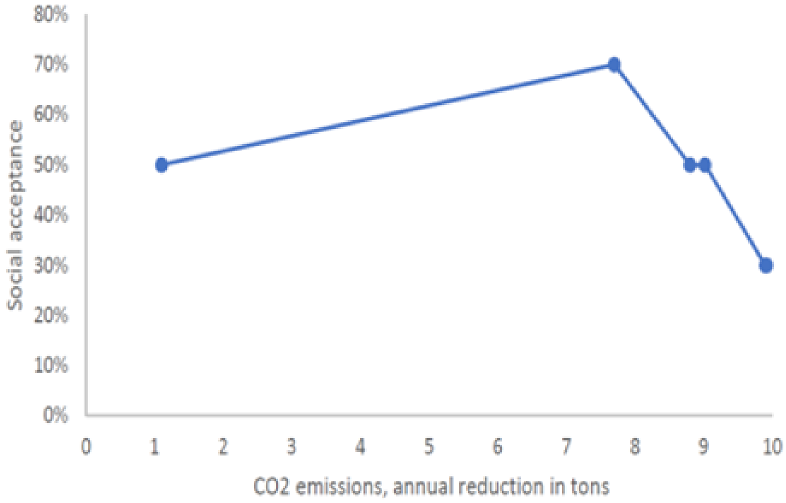

Fig. 6.3: Social acceptance vs. CO_2 reduction

Ultimately, the LASAGNE project confirms that engaging with communities from the project's inception and respecting their values and concerns transforms microgrids from mere technological infrastructure into shared projects based on cooperation and trust. The toolbox ensures that the shift from traditional consumption to cooperative prosumer roles is sustainable, both technically and socially.

The social acceptance toolbox developed in this chapter is not an isolated effort, but rather a foundational element that is interconnected with all the other components of the LASAGNE project. The success of technical innovations such as the Grid Edge Devices (GEDs), the flexibility applications, the forecasting models and the coordination mechanisms (Chapters 2-5) hinges on their alignment with user preferences and community values. By quantifying and optimising social preferences, the toolbox ensures that energy flexibility solutions are technically viable and socially embraced. This integration is also evident in the deployment strategies (Chapter 7) and future initiatives (Chapter 9), where insights into social acceptability directly inform the design and scaling of microgrid systems in various contexts. In essence, social acceptance acts as the connective tissue that binds the technological, economic and governance dimensions of LASAGNE together, enabling a truly cooperative and resilient energy transition.

6.8 References

- Fragniere, E., Sandoz, S., Abdenadher, N., Moussa, M., Serugendo, G. D. M., & Glass, P. (2023, June). Fostering “Energy Communities”: An Ethnographic-SECI Approach to User-Centered Residential Micro-Smart Grid Adoption. In 2023 11th International Conference on Smart Grid (icSmartGrid) (pp. 01-05). IEEE.

- Fragniere, E., Sandoz, S., Abdenadher, N., & Moresino, F. (2024, May). Developing card-based surveys: focus group insights for smart grid prosumer optimisation. In 2024 12th International Conference on Smart Grid (icSmartGrid) (pp. 579-581). IEEE.

- Fragniere, E., Sandoz, S., Abdenadher, N., Moussa, M., Serugendo, G. D. M., & Glass, P. (2023, May). Improving the social acceptability of microgrids. In 2023 IEEE Renewable Energy and Sustainable E-Mobility Conference (RESEM) (pp. 1-5). IEEE.

- Fragniere, E., Moresino, F., Sandoz, S., & Abdennadher, N. (2024, May). Enhancing Social Acceptance in Microgrid Implementation: A Share-of-Choice optimisation Approach. In International Conference on Informatics in Economy (pp. 289-299). Singapore: Springer Nature Singapore.

Chapter 7: Deployments

Authors: HES-SO//HEPIA (Raoul Dupuis, John White, Nabil Abdennadher), CLEMAP (Pascal Kienast), Giovanna Di Marzo (UNIGE), ElectriCITY (Josefin Danielsson)

7.1 Introduction

This chapter presents the deployment activities performed within the LASAGNE project. The objective is to demonstrate the operational integration of the developed framework, including hardware, software, and data-management components, in three pilot sites: two in Switzerland and one in Sweden. They provide complementary environments for validating technical performance, interoperability, and user acceptance of the LASAGNE framework.

7.2 Deployment Sites

In Switzerland, two pilots are considered: the school in *Les Vergers* eco-district in Meyrin and the cooperative housing site *CODHA / Rigaud* in Chêne-Bougeries. A third site, *Polygones* (Meyrin), has been prepared for subsequent extension.

Les Vergers constitutes an exemplary low-carbon district with high integration of renewable energy resources and local microgrid management. Seven *CLEMAP devices* were installed in the three buildings of the school. Fig. 7.1 shows an aerial view of the school (green rectangle) where we deploy the CLEMAP devices. They are installed across four electrical cabinets, selected from the twenty available throughout the school. These devices are used to either measure the overall consumption of the electrical panel or monitor the specific consumption of certain equipment connected to this electrical panel.

Fig. 7.1: Aerial view of Vergers school (46°13'54.6"N 6°04'07.4"E)

For the four selected electrical panels, a CLEMAP device was installed at the entry point to measure the total consumption of each panel. Additionally, CLEMAP devices were placed on certain circuits within the panels to obtain more granular measurements. Fig. 7.2 illustrates the installation of these devices along the four electrical panels.

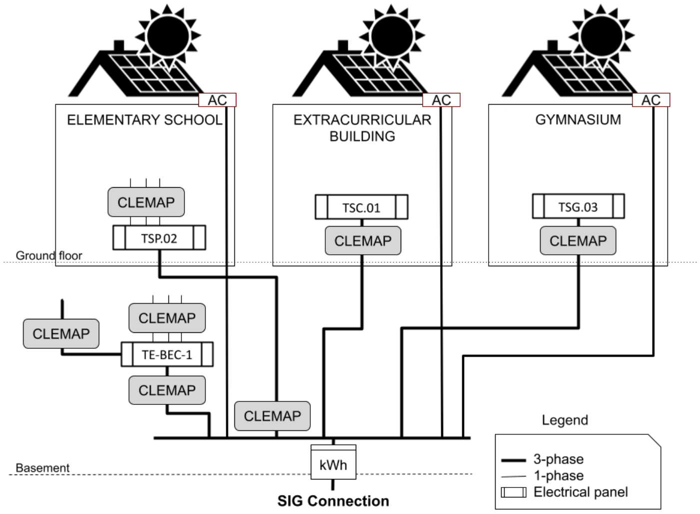

Fig. 7.2: CLEMAP devices in “Les Vergers” school

In the Chêne-Bougeries deployment (CODHA), we focus on a residential complex comprising six buildings (green rectangle in Fig. 7.3). Here, the CLEMAP devices have been strategically installed to monitor energy consumption in the common areas and the geothermal heating system. The selection of locations for these devices is guided by the objective of providing a comprehensive overview of energy use in shared spaces and specifically tracking the efficiency and consumption of the geothermal heating system, which is central to the complex's eco-friendly energy strategy.

Figure 7.3. Aerial view of CODHA (46°12'25.2"N 6°11'57.6"E)

With the installation of five CLEMAP devices, we aim to capture data that reflects the energy dynamics of the communal spaces and the heating system. This approach allows for a detailed analysis of energy consumption patterns and aids in identifying opportunities for energy conservation and optimisation. Each device is connected to the internet via cable, ensuring reliable data transmission for real-time monitoring and analysis. Fig. 7.4 presents a detailed view of the device installation locations; alongside the specific types of consumption they monitor. This visual representation aids in understanding the scope and focus of the energy monitoring efforts within the complex, highlighting our proactive measures in energy management and sustainability at Chêne-Bougeries (CODHA).

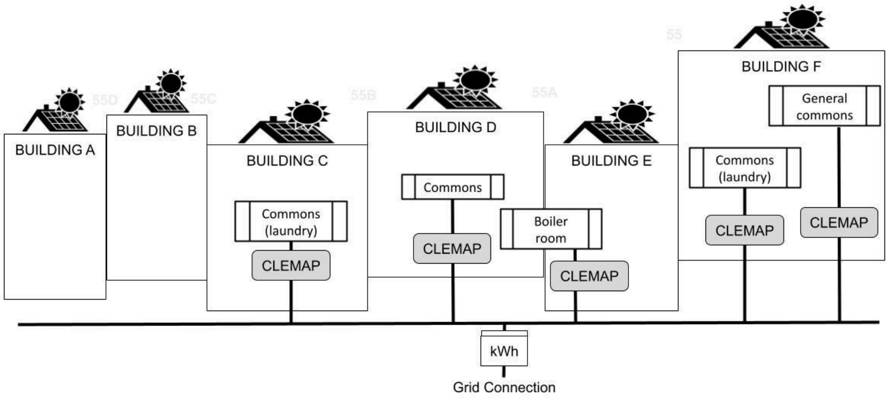

Fig. 7.4. Installation diagram of CLEMAP devices in Chêne-Bougeries

The Swedish pilot involves two housing associations, located in Hammarby Sjöstad, Stockholm. Since 2014, sustainability efforts in Hammarby Sjöstad have been driven by the citizen-led initiative ElectriCITY Innovation, a collaboration between residents, housing cooperatives, companies, public authorities, and academia. An important aspect of the work is to ensure that the innovations tested can be commercialized and further developed for international markets. The deployments focus on the integration of LASAGNE into existing infrastructures and evaluate applications such as dynamic tariffs, flexibility management, and peer-to-peer energy exchange.

Together, the Swiss and Swedish pilots cover a broad range of urban-residential conditions and stakeholder constellations.

7.3 Deployment Architecture and Integration

The Swiss deployment relies on *CLEMAP* devices based on Raspberry Pi hardware, configured to operate as *Grid Edge Devices* (GEDs). Each unit hosts lightweight micro-services responsible for local data acquisition and preliminary processing. The devices are integrated into the *Nuvla* platform, which provides comprehensive management of edge resources. Through *NuvlaEdge*, the consortium ensures remote device monitoring, software orchestration, service deployment, and automated updates. This architecture enables reproducible configuration across multiple pilot sites while maintaining security and scalability.

Resource constraints on the Raspberry Pi 3 platform have required targeted software optimisation to guarantee stable operation under continuous load. Key measures included reducing memory footprint through lightweight frameworks and optimised data-handling routines. In parallel, the consortium agreed on hardware upgrades to Raspberry Pi 4 to improve computational headroom, enable more demanding analytics at the edge, and support future LASAGNE applications requiring higher performance (see details in Chapter 2).

In Hammarby Sjöstad, pilot tests were carried out to evaluate the performance and interoperability of smart energy technologies within a real urban environment. The testing focused primarily on assessing the versatility and efficiency of battery systems in delivering key energy services, including peak shaving, spot price optimisation, reduction of grid-related costs, enhancement of self-consumption, and participation in Frequency Containment Reserve (FCR) and Fast Frequency Regulation (FFR) markets. These activities aimed to lower overall energy costs, generate new revenue streams, and strengthen grid stability, thereby contributing to the district’s broader ambition of establishing a resilient and sustainable urban energy infrastructure. Further, tests were conducted to explore the feasibility of enabling participation in frequency-controlled ancillary service markets using electric vehicle (EV) chargers from the brand ChargeAmp. These tests were performed under real-life conditions and included an evaluation of the technical performance and controllability of the EV chargers through the Cleamp platform. Together, these experiments provided valuable insights into the potential of integrating distributed energy resources for enhanced flexibility and grid support within a local energy community context.

7.4 Preliminary Results and Insights

The deployment activities confirm the operational feasibility of the LASAGNE concept. In Switzerland, GEDs operated reliably under continuous load with stable communication. In Sweden, the integration with legacy systems proceeded without major compatibility issues, and early tests of dynamic-tariff mechanisms produced consistent responses from connected devices.

Initial quantitative results include improved RAM memory efficiency, average network latency below 150 ms, and uptime exceeding 98 %. Qualitative feedback from user studies indicates positive attitudes toward local energy sharing, combined with sensitivity regarding data privacy and tariff transparency.

The pilot in Hammarby Sjöstad confirms that battery storage, when integrated with a marketplace‑oriented EMS, can stack local and system‑level services without compromising residential comfort, making a solid business case for housing associations. The stringent Nordic pre‑qualification results demonstrate technical readiness of mid‑scale BESS for FCR/FFR and validate LASAGNE’s coordination‑model abstraction for real‑time controls. Conversely, the EV‑charger study reveals that cloud‑mediated APIs introduce unacceptable latency for frequency services. Direct OCPP control and tighter time‑synchronization are mandatory. This insight highlights the value of negative results and the need for iterative hardware-in-the-loop testing early in product lifecycles.

The Hammarby Sjöstad pilot has provided valuable insights into the integration of flexible energy assets in housing associations within energy communities, highlighting both the potential and the complexity of broader adoption. It highlights the potential for increased energy efficiency, cost savings, and local flexibility, while also showing the importance of clear governance and collaboration. Since many housing associations are managed by resident boards without technical expertise, supportive frameworks and transparent role allocation between technology providers, service companies, and associations are essential to enable confident decision-making. The tests also underline the value of developing integrated control solutions that connect different energy systems and suppliers, ensuring coherent operation and performance. Tailored business models and shared value frameworks can further strengthen engagement and scalability. With clear economic incentives and trust-based governance, housing associations can play an active role in the evolving landscape of energy communities and contribute to a more resilient and sustainable energy system.

7.5 Challenges and Mitigation

1. Several challenges were encountered during deployment.
   Technical constraints: limited processing capacity at the edge required optimisation and modularisation of services.

2. **Connectivity issues:** existing building networks occasionally lacked sufficient coverage, necessitating dedicated communication gateways.

3. **Data protection:** user concerns about data confidentiality were addressed through anonymisation, encrypted transmission, and clear consent procedures.

4. **Stakeholder coordination:** multi-actor environments (housing cooperatives, municipalities, utilities) required structured communication and defined responsibilities.

The Hammarby Sjöstad pilot highlighted challenges in integrating flexible energy assets within housing associations, along with ways to address them. Governance and decision-making can be complex for resident-managed boards, but clear role definitions, guidance, and support can build confidence. Coordination across multiple energy systems requires integrated management platforms and clear operational responsibility. Tailored business models and demonstrable economic benefits, such as cost savings or new revenue streams, are essential to enable adoption. Together, these measures can help housing associations actively contribute to resilient and sustainable energy communities.

These mitigation measures ensured stable operation and acceptance of the deployed systems.

7.6 The marketplace platform

This section deals with the marketplace concept developed within LASAGNE. The objective of this marketplace is to define and establish the contractual relationships that govern the financial flow between the five parties involved: Need Owners (end-customers households and microgrid owners, communes, municipalities), the systems integrators (DSOs, LECs owners/admin, etc.), the independent software vendors (Flexibility applications developers), the edge equipment vendors and the marketplace platform provider.

When the project was submitted, it was planned that LASAGNE would use SixSq's Nuvla.io platform to develop a marketplace tailored to the electrical energy sector. At the end of 2021, SixSq was acquired by the Ekinops group, and in 2024, Ekinops decided not to extend Nuvla.io's Marketplace into a model supporting the needs of the project.

Therefore, we will simply present here the marketplace concept (No implementation has been conducted): how it applies to edge and cloud architectures in general and how this concept will be leveraged in the energy sector by the LASAGNE project. To succeed, edge-based business must involve a number of stakeholders (see Fig. 7.5):

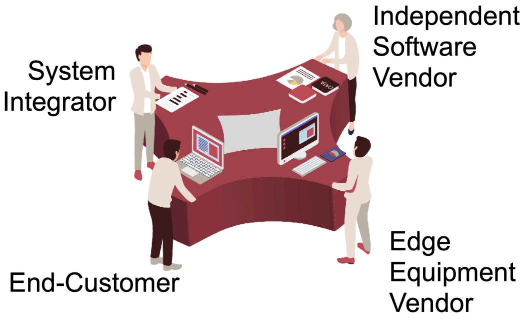

Fig. 7.5: Stakeholders for edge-to-cloud ecosystem

First, we have an **End-Customer** (Need owners) who wants to make sense out of the data it has, or manage its connected sensors/devices/machines/things, all this to understand its environment, and all connected assets that need to be monitored and controlled.

Second, we often have a **System Integrator** (SI) working for our end-customer to deploy and maintain an edge infrastructure and services, such that sensors and actuators can be controlled locally and just enough processed data transferred to the cloud or data center.

Third we have **Independent Software Vendors** (ISV) providing the applications required to run at the edge, that once deployed process sensor data in real time and turn this data into business relevant information and actions. These applications are often AI or Machine Learning apps.

And finally, we have **Edge Equipment Vendors** (EEV) providing the edge hardware to which all sensors are connected.

In LASAGNE, the marketplace was proposed to provide the software infrastructure required to bring all the stakeholders together, for example:

- End-Customers: households but also microgrid owners, communes, municipalities
- SI: power electric companies such as DSOs
- ISV: Tvinn/RECAP, CLEMAP
- Edge Equipment Vendor: CLEMAP device

An attractive aspect of the marketplace is its ability to establish a simple digital contractual framework, such that purchases and fund distribution is automated, simple and secure.

For example, the marketplace platform should allow an ISV to set the price it wants for each app registered in the app store. The ISV can also either accept to follow the standard terms set by the platform, or bring its own terms. Then, when a SI deploys an app on its customers' edge infrastructure, it accepts the terms and price, and the platform executes the app deployment. Since all stakeholders have accepted the terms of the platform, it can then issue invoices and payment instructions. Upon reception of payments, the platform then redistributes the payments to the ISV (typically 80% of the app store price) and the platform operator.

With this level of automation in place, ISVs can be added to the platform, and SIs can choose which app and ISV to use. In agreement with the regulation and End-Customers, the SIs put in place a governance policy, such that only apps having been certified are allowed to be deployed.

7.7 Conclusion

This chapter summarises the deployment and validation of the LASAGNE framework in Switzerland and Sweden.

In Switzerland, the initial deployments used the Raspberry Pi 3 version of CLEMAP. This version does not support the features required by the project. However, it did enable us to collect the data needed to develop predictive energy consumption models. Thanks to the LASAGNE project, a Pi4-based version was tested and implemented. It is currently being deployed in Meyrin (Polygones) and other sites in Switzerland (See details in Chapter 9).

In Sweden, pilot activities in Hammarby Sjöstad focused on integrating LASAGNE with existing energy infrastructures. Tests of battery storage and EV charging demonstrated the potential for flexibility services, market participation, and grid support, though latency in cloud-based EV control remains a limitation to be further explored. Overall, the deployments validate the technical feasibility of LASAGNE's edge-to-cloud architecture and highlight governance, interoperability, and optimisation as key enablers for scalable, community-based energy systems. The chapter concludes by presenting the concept of a marketplace designed as part of the project.

Chapter 8: Dissemination activities

This chapter list the dissemination activities carried out within LASAGNE during 2024 and 2025:

1. Towards a self-adaptive Federated Learning (FL) based edge-to-edge framework. Application to energy sector (keynote), AROSA workshop, ECSA'2025 conference, Limasol, Cyprus, Sept. 15 – 19, 2025.
2. 13 June 2025: Tutorial presenting the results of LASAGNE (distributed Federated Learning): see this link
3. Public event organised by the KTH and Electriciyin Stockholm, 22 May 2025. See this link
4. A ML-based edge-to-cloud platform for digital energy services, IEEE 6th International Conference on Smart Information Systems and Technologies (SIST'2025) (keynote), Astana, Kazakhstan, May 13-15, 2025
5. A ML-based edge-to-cloud platform for digital energy services, "Forum international de la transition énergétique", Béconcour and Drummonville (Quebec, Canada), 20-23 January 2025
6. Towards a Continuum Computing framework for decentralised Federated Learning Based Context-aware IoT Applications (keynote), GCET'2024, Gran Canaria, Spain, December 9-11, 2024
7. Towards a decentralised Federated Learning based edge-to-edge Continuum Computing framework (tutorial), IC2E'2024, Paphos, Cyprus, September 24-27, 2024.
8. A distributed Federated Learning approach for a Continuum Computing platform (keynote), WORIE'2024, Ljubljana, Slovenia, June 17-18, 2024
9. Towards a distributed Continuum Computing platform for Federated Learning Based Self-Adaptive IoT Applications (keynote, MECO'2024 and CPSIoT'2024, Budva, irnegro, June 11-14, 2024.
10. Presentation at Unige Data Science Day - Giovanna Di Marzo Serugendo, "*Multi-agents systems - a tool for modelling and developing complex systems*" - 15.09.2022 https://datascience.unige.ch/en/research/uniges-data-science-days/uniges-data-science-days-3
11. Poster presentation at Unige Data Science Day - Philippe Glass and Giovanna Di Marzo Serugendo, Computer Science Centre, "*Lasagne - implementation of smartgrid in the form of a coordination system of digital twins*" - https://archive-ouverte.unige.ch/unige:164521 15.09.2022

https://datascience.unige.ch/en/research/uniges-data-science-days/uniges-data-science-days-3

12. 34th Séminaire CH «Gouverner aujourd'hui: La Suisse en 2040» - Presentation to all Swiss Cantonal ministers organised by Fondation pour la collaboration confédérale. Giovanna Di Marzo Serugendo, Informer de manière efficace les politiques publiques grâce aux technologies numériques, l'IA et les jumeaux numériques - 05.01.2024
13. Keynote speaker. Knowledge Management Conference 2024. Giovanna Di Marzo Serugendo, Effectively Informing Public Policies with Digital twins, AI and digital technologies – 04.07.2024 https://www.iiakm.org/conference/KM2024/pdfs/KM2024_KEYNOTE_DiMarzoSerugendo.pdf
14. LASAGNE projects and results proposed during continuous education programs on AI at Unige (from 2023 to 2025)
15. Conférence Mairie de Gaillard - Partage des Savoirs - Giovanna Di Marzo Serugendo, Viola Krebs, Lamia Friha. L'Intelligence Artificielle en Action - Transformations Sociétales et Défis - 4.11.2025. Press release:
    https://www.ledauphine.com/culture-loisirs/2025/10/30/une-conference-pour-tout-savoir-sur-l-intelligence-artificielle-le-mardi-4-novembre
    https://www.francebleu.fr/emissions/l-invite-a-la-une-ici-pays-de-savoie/giovanna-di-marzo-serugendo-l-ia-n-est-ni-une-menace-ni-une-solution-miracle-4656743
    https://www.lemessager.fr/649343461/article/2025-10-28/chatgpt-bouleverse-notre-rapport-la-technologie-une-professeur-nous-dit-tout-sur

Chapter 9: Beyond LASAGNE

Authors: HES-SO//HEPIA (Raoul Dupuis, John White, Nabil Abdennadher), KTH (Carlo Fischione), CLEMAP (Pascal Kienast), Giovanna Di Marzo (UNIGE)

9.1 Introduction

The goal of the LASAGNE project was to develop and deploy, as Proof of Concept, a set of forecasting models for energy consumption (LASAGNE framework) that will be used in different contexts. We were particularly interested in two contexts: households and EV charging stations. However, the concept can be easily extended to other contexts such as microgrid (LEC), buildings, etc. The LASAGNE framework is assumed to be the “building block” used to develop and deploy flexibility applications (grey boxes in Fig.1.1)

LASAGNE has also implemented a set of tools (ML pipeline for energy forecasting and visualisation tools) to enable:

1. Software vendors to automate the ML pipeline (Data cleaning, Feature extracting and training process)
2. End-users to intuitively visualise prediction results and actual data.

The prediction models were deployed on edge-to-cloud solutions, enabling compliance with data privacy regulations, reducing cloud congestion, and optimising communication between IoT devices (that measure electricity consumption/production) and the cloud.

LASAGNE also developed and tested a decentralised Federated Learning approach to improve the accuracy of the prediction models. The idea is to deploy (in the future) this approach on edge-to-edge architectures.

For the flexibility applications to spread out, they must be accepted and used by the different stakeholders. Whether these actors would tolerate such services is far from obvious and acceptance constitutes a potential barrier for the adoption and diffusion of flexibility applications. LASAGNE tackled this problem by proposing an approach to include this aspect from the beginning. This methodology demonstrates how to construct a progressive toolbox for achieving 'social acceptance by design' within smart micro-grids, moving systematically from qualitative community research to quantitative design optimisation.

The expertise and results acquired during the LASAGNE project have enabled the consortium to initiate new projects, which are detailed in this chapter.

The chapter is organised as follows: the first section deals with three applied research projects launched thanks to the LASAGNE projects: BATTwin, Smart Energy District (SED) and O-CEI.

KTH is involved in the BATTwin project while HES-SO and CLEMAP are involved in the two projects SED and O-CEI. Thanks to the support of the commune of Meyrin (Geneva) who signed a support letter for LASAGNE, the Polygones association is also a partner of the O-CEI project.

The second section of this document presents the projects initiated by CLEMAP thanks to the results of LASAGNE

9.2 Research projects

9.2.1 BATTwin

LASAGNE's results have enabled KTH to take part in the European BATTWIN project, which started in January 2024. Li-ion batteries are fundamental components for the energy transition of the European eco-system. Currently Europe lags behind Asia in terms of Li-ion battery cell manufacturing and more than 90% of the world's production takes place in China, Korea and Japan. To overcome this situation, there is an ambitious ramp-up plan of 25 new gigafactories in Europe with an expected value of €35 billion by 2030. However, in the ramp-up phase of these Gigafactories, a massive production of defects is expected, between 15% - 30%. The new European Gigafactories will also bring demand for €150 bn of battery manufacturing equipment. To support this demand, the EU production equipment industry needs to fill the current knowledge gap and gain competitiveness towards Asian providers, grounding on its world-wide leadership in high-tech, green technologies, enhanced by industry 4.0 digital solutions, exploiting the European Zero-Defect Manufacturing paradigm. The objective of BATTwin is to support this scenario by developing a novel Multilevel Digital Twin platform towards Zero-Defect Manufacturing in battery production, that will reduce defect rates in battery production lines.

9.2.2 Smart Energy District (SED)

The SED project[17] (2024 - 2027) is one of the HES-SO flagship projects (1.5 M CHF). SED will create a Smart Energy District where energy consumption information can be accessed and flexible components can be controlled to optimally manage the distribution network and the environment offers a high degree of standardization and intercompatibility, allowing various stakeholders to deploy their services using the same infrastructure and combining sectors.

A digital platform will be developed for operating the Smart Energy District to control flexible resources in real time, ensure network operational security, and provide a toolbox for future digital services applied to energy. The project consortium, in close collaboration with Groupe-E, has selected the Marly Innovation Center[18] as the first site to deploy this platform.

The Marly Innovation Center (MIC), located in the city of Marly (Fig. 9.1), in the greater Fribourg area is one of Switzerland's largest technological campuses. Its unique ecosystem

---

[17] https://www.hes-so.ch/smart-energy-district
[18] https://marly-innovation-center.org/en/accueil-en/

offers its tenants a wide range of facilities and infrastructure, such as laboratories, offices, production space and storage facilities, spread over 70000 m2 of rental space, among which 7000 m2 of laboratories. Today, the site hosts over 180 companies of all sizes, representing a workforce of more than 800 employees. In addition to state-of-the art facilities, the recently-developed EcoQuartier (first of its kind in the canton of Fribourg), a SEED-certified residential complex, will be able to house more than 2500 residents in the coming years.

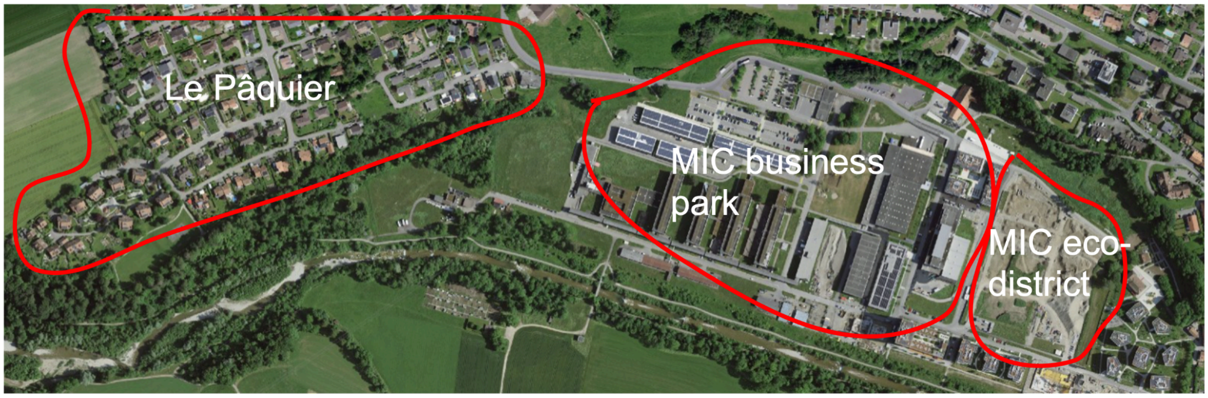

Fig. 9.1 : Marly, Fribourg, CH

The district's extensive infrastructure facilitates the expeditious initiation of energy control and demand response measures that do not necessitate new construction. The existing installations, which include solar photovoltaic (PV) systems, heat pumps (HPs), and interconnected thermal networks, offer considerable potential for integrating electrical systems and enhancing grid stability via thermal network flexibility.

The machine learning experience of the LASAGNE project will be directly applied to the Smart Energy District. The energy consumption data currently collected from the LASAGNE sites will also be added to the Smart Energy District data.

9.2.3 O-CEI

The O-CEI project[19] will use some of the sites of the LASAGNE project to run a Swiss pilot to measure the social engagement and acceptability of energy flexibility in urban areas.

The goal is to foster prosumers (energy consumer/producer) into an increasingly liberalized energy market, enabling them to reap new forms of electricity production and consumption. This pilot will be deployed at three cantons in Switzerland: Geneva, Valais and Fribourg (to be confirmed), that engage users and energy flexibility in three types of buildings: urban apartments, administrative offices and tourism (skiing) infrastructure.

---

[19] https://o-cei.eu/

Prosumers are called to play a prominent role in the energy transition to a renewable energy mix. Here, engagement requires adoption of new technologies and social acceptance.

The Swiss pilot will demonstrate how innovative Smart IoT solutions (already) in place can become part of an integral (O-CEI) platform, increasing such acceptance and user engagement while deploying Smart Energy Services promoted by advanced Cloud-Edge-IoT utilities and technologies.

Geneva (Polygones, Meyrin school and Chêne-Bougeries,Codha)

The Polygones (Fig. 9.2) deployment will be enhanced through the deployment of CLEMAP energy measuring devices alongside the current eSmart devices. The users that allow their energy usage to be monitored, will be able to view their energy usage through a portal.

Fig. 9.2 : Polygones, Meyrin, Geneva, CH

Meyrin school and Codha deployments are described in the LASAGNE progress reports 1 and 5. There are 7 (*resp*. 5) CLEMAP devices already deployed in the Vergers school (*resp*. Codha buildings in Chêne Bougeries). These deployments will be extended within the O-CEI project.

Valais: The Val d'anniviers is an Alpine valley home to numerous villages including the resorts of St-Luc, Chandolin, Grimentz, Zinal and Vercorin. A distributed municipality of permanent residents and tourists within a complex of 700 km of hiking and biking trail and

220 km of ski pistes. In and around the valley of Val D'Anniviers, edge equipment (e.g., sensors, actuators, energy meters and computing elements) will leverage the O-CEI orchestration engine, monitoring and cybersecurity to enhance the previously existing CEI solution.

Confirmed installations include the Boutique Hotel Chandolin (Fig. 9.3) and the ski lifts of St Luc-Chandolin. In the above hotel, CLEMAP devices with 80 Amps capacity will be deployed to monitor the energy consumption of: electric vehicle chargers; hotel laundry; sauna-spa; guest rooms; restaurant kitchen. In the case of the ski lift company of St Luc-Chandolin, there will be three ski lifts, one restaurant (at altitude) and one funicular railway equipped. The expected power consumption for these installations ranges from 400 to 2000 Amps.

Fig. 9.3 : Boutique Hotel, Chandolin, Valais, CH

9.3 Industrial projects

The innovative solutions in the field of load and energy management for electric vehicle (EV) charging play a significant role in enhancing energy efficiency and promoting the integration of renewable energy sources. In Switzerland and Europe, thanks to our work within LASAGNE, CLEMAP is able to commercialize forecast-based applications that adjust charging processes in real time based on grid demands and market conditions. This not only helps stabilize power grids but also enables users to benefit from more cost-effective energy pricing models and maximize the use of sustainable energy sources, particularly solar (PV).

The applications are designed for a wide range of customers, including individual EV owners, fleet operators, energy utilities, charging infrastructure providers, and businesses with large-scale charging needs. Whether it's residential users looking to optimize their EV charging at home or commercial fleets seeking to manage large-scale charging operations efficiently, our solutions provide the flexibility and intelligence required to integrate EVs into a sustainable energy ecosystem.

Additionally, we are expanding the use of our technology to industrial-size batteries, where the same intelligent load management systems will be deployed to optimize energy usage, maximize the use of renewable energy, and support grid stabilization. Energy storage operators can leverage our technology to ensure that stored energy is used efficiently during high-demand periods or when renewable energy is abundant, further enhancing the overall energy management strategy.

Energy utilities can leverage our technology to better manage grid stability and demand-response strategies, while charging station operators can enhance their service offerings by providing dynamic pricing models and optimizing energy usage. Solar energy producers and storage system operators can also benefit by ensuring that their excess renewable energy is effectively utilized for EV charging, maximizing their return on energy assets.

The applications seamlessly integrate with existing and deployed energy management systems, providing both consumers and businesses the opportunity to actively contribute to the energy transition through flexible loads, optimized use of renewable energy, and the integration of energy storage solutions.

Here a few projects that have grown out of LASAGNE and where CLEMAP will further deepen and specify the results:

Laderabatt (SFOE funded project), between 2024 and 2026: As part of the *Laderabatt* project, CLEMAP's current technical infrastructure, including its load management system, is being used to enable operators of charging points at public parking lots, as well as employee, visitor, and private parking spaces, to offer discounts for charging. Electric vehicles are an important alternative to fossil-fuel-powered cars, but costs are a key decision factor in times of rising electricity prices. The project aims to counteract this.

By aggregating charging points into a pool (virtual power plant), which is connected to the secundary balancing energy market via BKW Energie AG as flexibility provider for flexibility purposes, charging point operators (and finally users) are compensated for the available flexibility and receive remuneration. EV users charging at their destination can simply plug in their cars and will only marginally notice any charing delays while potentially benefit from discounted rates if electricity needs to be drawn at that moment for grid-supportive purposes.

LASAGNE's forecasting model would serve here to indicate the charing power which is forecasted to the pooling provider. The project is supported by important CLEMAP's customers such as Ottofischer, Energie360, EWZ and partners, webinars and stakeholder events demonstrated that charge point operator see a new business model in the approach. CLEMAP aims to test both market acceptance, technical feasibility and the business model while hopefully commercializing the application by March 2026.

Chargesync+ (SFOE funded project), between 2025 and 2028: The project goal is to use dynamic electricity tariffs as incentives for end users with electric vehicles at home and at the workplace, enabling them to provide the flexibility of their charging processes for grid-supportive load management via a user interface. The Sintio mobile app serves as the user interface, allowing end users to easily indicate their willingness to respond to dynamic electricity tariffs.

CLEMAP’s manufacturer-independent load management system translates this user feedback into a smart charging plan. Sintio’s billing solution ensures user-friendly invoicing based on 15-minute load profiles from the charging stations, read out via OCPP.

LASAGNE’s forecasting models shall help to improve the charging algorithm considering not only the tariff but as well the time to reach a respective state of charge for an EV.

In field tests with the energy providers ewl (Lucerne), ESB (Biel), EnBAG (Brig), the municipal utility of Winterthur, and the property management company Seeblick Immobilien (Bäch SZ), the practical feasibility and market viability of the approach are being tested. The project will be terminated in 2028.