2026 22nd International Conference on the European Energy Market (EEM) | 979-8-3195-3554-2/26/$31.00 ©2026 IEEE | DOI: 10.1109/EEM68581.2026.11589597

Scientific Machine Learning in Local Energy Systems

Samrat Bose[1,2], Jaisiel Santana[2], Peter Tzscheutschler[1], Reinaldo Tonkoski[1]
[1]Technical University of Munich, Arcisstraße 21, 80333 Munich, Germany
[2]European Institute for Energy Research, Emmy-Noether Straße 11, 76131, Karlsruhe, Germany
samrat.bose@tum.de, samrat.bose@eifer.org

***Abstract*—The focus on the development of control strategies for Local Energy Systems is a hot topic of research for the inclusion of distributed renewable energy resources into the energy ecosystem. Although there has been significant research on the application of pure data-driven Machine Learning methods to develop these control strategies, the paradigm of combining scientific computing with the data-driven approaches through a multi-model method is still unexplored. This paper provides one of the first multi-model approaches of scientific machine learning in a real demonstrator site. The architecture proposed connects Agent-Based Modeling(Anylogic-Java) for stakeholder simulation with Pandapower(Python) for power flow analysis, facilitated by state-of-the-art transformer-based forecasting (Chronos 2) and Model Predictive Control (MPC)(Python). The results demonstrate a reduction of peak load and equivalent battery cycles and better performance in comparison to pure data-driven forecasting.**

***Index Terms*—Energy Communities, Scientific Machine Learning(SciML), Physics Informed Neural Networks (PINN), Multi-model methods, Model Predictive Control**

I. Introduction

The global energy transition with inclusion of distributed renewable sources at large scale is redefining the electricity networks. Inclusion of energy communities (RECs) by EU Commission to achieve carbon neutrality by 2050 [1] through clean energy package is demonstrating enhanced penetration of distributed renewable energy resources (DER) in the distribution grids. However, this is also causing significant operational challenges for the Distribution System operators (DSOs) in terms of voltage fluctuations, line congestion, and reverse power flows [2].

Local Energy Communities (LECs) have emerged as one of the key players for democratization of energy access and inclusion of DERs. The application of various strategies to implement local self-consumption, peer-to-peer trading, and other strategies, the LECs can provide better independency from the grid, reduce infrastructure investments for capacity expansion, and in turn provide flexibility services to the DSOs [5]. However, such inclusion requires advanced control strategies which can mitigate the uncertainty of renewable generation, the mismatch of supply and demand, and the physical constraints of the grid.

The inclusion of battery energy storage system (BESS) in several LECs by implementing different methods of demand response (DR) makes it invaluable for achieving flexibility and provide grid support services [3]. However, the control strategy deployed for the operation of the BESS is critical for its effectiveness. Traditional heuristic controllers, although simple and robust, are reactive in nature and fail to capture the futuristic outlook in its strategy. Hence, the paradigm of predictive control offers a better paradigm for optimization over a time period while also including the system constraints.

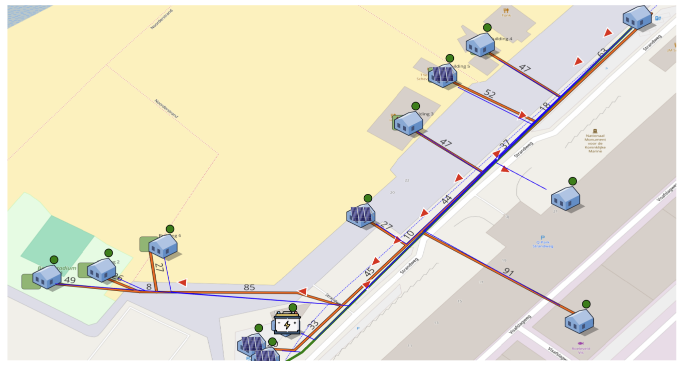

Fig. 1. Visual Representation of power flow of a Local Energy Community in a pilot site in U2Demo Horizon EU project [7] [8] [9]

In order to develop such sophisticated strategies, the advancement of research and application in the field of Machine Learning(ML), and more specifically the combination of scientific computing with data driven machine learning presents an interesting paradigm for exploitation and implementation. For e.g. Amazon's Chronos 2, a pre-trained transformer model, achieves state-of-the-art performance on energy forecasting benchmarks while providing well-calibrated predictions [6].

Despite these technological advances in various fronts, significant gaps remain in the deployment of such advanced tools and technologies. Most studies involving ML based energy management either rely on simplified simulation models which ignore power flow, whereas the power system studies often intend to use pure data driven forecasts or other simplified statistical models. Another significant challenge revolves around the interoperability of such tools which inhibits the heterogeneous simulation frameworks, that may capture the ML frameworks, agent-based models, and power flow solvers in one co-simulation environment. This paper addresses these gaps by presenting a multi-model co-simulation framework that integrates Chronos 2 for forecasting, Model Predictive Control with physics constraints for optimal battery dispatch, Agent-based modeling in AnyLogic for representing community stakeholders and their interactions, and PandaPower for accurate AC power flow analysis in a real demo site pilot under the U2Demo Horizon EU project as demonstrated in figure 1.

The paper is organized as follows. Section II reviews relevant literature on scientific machine learning in energy systems and more particularly on local energy communities. Section III describes the model methodology, Section IV denotes the simulation setup, and the Dutch pilot site configuration. Section V presents preliminary results and discussion. Section VI concludes with key findings.

II. Application of SciML in LEMs

A. Scientific Machine Learning

One of the earliest well established definition of Scientific Machine Learning (SciML) comes from [17] which states that scientific machine learning integrates data-driven machine learning with physical laws, differential equations, or domain models to improve prediction, interpretation, and generalization. SciML has found its various application in energy sectors in the last couple of years. The application of Physics Informed Neural Networks (PINN) is demonstrated for various purposes in the last couple of years like energy management in [18], improving building efficiency in [19], and solving non linear energy supply-demand systems [20]. Reinforcement Learning (RL) with some physical constraints has been demonstrated as a novel solution in recent research where it found its application in various sectors [22], particularly in the field of energy management systems [21], control strategy for energy efficient buildings [23], and energy community management [24].

B. SciML in Local Energy Systems

Local Energy systems comprises of different topics of research including but not limited to local energy markets, energy communities, low voltage grid operations, etc. The application of machine learning for local energy communities and local energy markets for data generation, operation and control strategies is an evolving topic of current research [13] [14], [15]. The application of machine learning techniques for energy trading is demonstrated in [25] [26]. The application of PINN in local energy systems is demonstrated in some research articles for surrogate model for grid dynamics in [27], in urban building energy modeling [28], in real-time grid-forming inverter control in local energy systems in [29], and in forecasting of PV for local energy systems for energy management [30].

C. Research Gap

Despite the evolving research in the field of SciML for energy sector, there remains critical gaps in its application for local energy systems. First of all, existing research on foundational AI models on forecasting focuses on accuracy benchmarks rather than operational integration in real-time control systems. Secondly, model predictive control studies primarily focuses on synthetic or simulated forecasts rather than advanced ML predictions undermining the value of state-of-the-art forecasting methods. Thirdly, co-simulation frameworks for energy communities rarely integrate all three layers: ML forecasting, physics-based optimization, and agent-based behavioral modeling. This paper makes the following contributions to address these gaps as demonstrated in figure 2:

1) First Real-World Demonstration: We present the first implementation of a transformer-based foundation model (Chronos 2) for real-time control of a local energy community with actual meter data. 2) Multi-Model Co-Simulation Framework: We develop a lightweight integrated architecture enabling real-time data exchange between AnyLogic (agent-based simulation), Python ML in a server (Chronos 2 forecasting), and PandaPower (AC power flow) without requiring complex middle-ware standards. 3) Physics-Constrained SciML: We demonstrate how ML forecasts can be integrated with MPC optimization while maintaining power flow constraints, ensuring that control actions are both data-informed and physically feasible. 4) Comparative Evaluation: We evaluate three control strategies (rule-based, ML Heuristic, SciML+MPC) using realistic KPIs including battery cycles, transformer peak loading, and computational latency on a real pilot site.

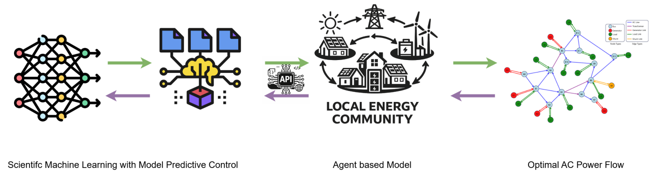

Fig. 2. Multi-model framework for application of SciML in LEC [7] [8] [11]

III. Model Methodology

A. System Architecture

The architecture comprises of a co-simulation framework facilitating integration of Agent-Based Modeling (ABM) with physics-based power flow analysis and Scientific Machine learning. The model consists of three primary layers: (1) the *Agent Layer* implemented in AnyLogic, representing grid components as autonomous agent populations; (2) the *Power Flow Layer* utilizing PandaPower for AC power flow analysis; and (3) the *Intelligent Control Layer* comprising the Machine Learning (ML) forecasting service and Model Predictive Control (MPC).

The co-simulation follows a time horizon of 2 weeks with a discrete time step Δt. At each timestep t, the ABM layer updates agent states, aggregates net power injections $\mathbf{P}_t^{inj}$, and exports them to PandaPower via a `PyCommunicator` interface. The Newton-Raphson solver returns nodal voltage magnitudes $V_{i,t}$ and line loadings $L_{ij,t}$, which are used for constraint evaluation and agent state updates and sent to the Machine learning model for prediction following the constraints.

B. Network Modeling

The distribution network is modeled as a directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N}$ represents the set of buses and $\mathcal{E}$ the set of lines. Each bus $i \in \mathcal{N}$ is characterized by nominal voltage V_{nom}, bus type $\tau_i \in \{slack, PQ\}$, and net complex power injection $S_{i,t}^{inj} = P_{i,t}^{inj} + jQ_{i,t}^{inj}$ [MVA]. The slack bus maintains $V_0 = 1.02$ p.u. For PQ buses, the net active power injection is:

$$P_{i,t}^{inj} = \frac{1}{1000}\left(P_{i,t}^{gen} - P_{i,t}^{load} + P_{i,t}^{batt}\right) \quad \text{[MW]} \tag{1}$$

where $P_{i,t}^{batt}$ is positive for discharge and negative for charge. Reactive power is calculated assuming a constant power factor pf. Lines $(i,j) \in \mathcal{E}$ are modeled using series impedance $Z_{ij} = R_{ij} + jX_{ij}$ derived from cable parameters:

$$R_{ij} = \frac{r_{\Omega/km} \cdot l_{ij}}{n_{parallel}}, \quad X_{ij} = \frac{x_{\Omega/km} \cdot l_{ij}}{n_{parallel}} \tag{2}$$

The AC power flow problem is solved using PandaPower's implementation of the Newton-Raphson method. The solver determines nodal voltages $\mathbf{V}_t$ and line currents $\mathbf{I}_t$ that satisfy the power balance equations:

$$S_{i,t}^{inj} = V_{i,t} \sum_{j \in \mathcal{N}} Y_{ij}^* V_{j,t}^* \tag{3}$$

where $\mathbf{Y}_{bus}$ is the nodal admittance matrix. Line loading percentages are calculated as:

$$L_{ij,t} = \frac{|I_{ij,t}|}{I_{ij}^{max}} \times 100\% \tag{4}$$

C. Integration with Machine Learning Models

1) Data Exchange Interface: The co-simulation interfaces with the ML forecasting service through HTTP REST calls at every simulation step. AnyLogic exports historical load and generation measurements as CSV files (one per asset category), each containing an ISO-8601 timestamp column followed by power measurements in kW. A POST request transmits data with a minimum look-back window of $H = 1344$ steps (14 days at 15-min resolution). The service returns an aggregated CSV covering a 24-hour horizon ($K = 96$ steps) with columns for timestamp, category, and quantile forecasts (e.g., `_q10`, `_q90`). The point forecast $\hat{P}$ is computed as the mean over sampled trajectories.

2) Physics-Informed Constraint Integration: To ensure physical feasibility of ML forecasts without degrading model accuracy, a post-processing constraint layer (Strategy A from III-D2) is applied per asset:

$$\hat{x}_{T+h}^{valid} = \text{clip}\left(\hat{x}_{T+h},\ x_i^{min},\ x_i^{max}\right), \tag{5}$$

guaranteeing hard constraint satisfaction (e.g., non-negative power, contracted limits).

D. Control Strategies

The BESS agent maintains a state-of-charge SOC_t governed by the discrete-time dynamics:

$$SOC_{t+1} = SOC_t + \frac{\eta_{rt} \cdot P_t^{ch} \cdot \Delta t - \frac{P_t^{dis}}{\eta_{rt}} \cdot \Delta t}{E_{cap}} \tag{6}$$

where E_{cap} is the usable energy capacity, η_{rt} is the round-trip efficiency, and P_t^{ch}, P_t^{dis} are the charging and discharging powers, respectively. The battery operates under the following constraints:

$$0 \le P_t^{ch} \le \min\left(P^{max}, \frac{E_{cap}(SOC^{max} - SOC_t)}{\eta_{rt}\Delta t}\right) \tag{7}$$

$$0 \le P_t^{dis} \le \min\left(P^{max}, \frac{E_{cap}(SOC_t - SOC^{min})\eta_{rt}}{\Delta t}\right) \tag{8}$$

$$SOC^{min} \le SOC_t \le SOC^{max} \tag{9}$$

Three hierarchical control modes are implemented: (1) *Rule-Based Control* (RBC), using real-time net load thresholds; (2) *ML Heuristic Control*, applying predictive threshold logic using 24-hour ahead forecasts $\hat{\mathbf{P}}^{net}$ to modify charging/discharging decisions; and (3) *SciML MPC Control*.

The MPC layer solves a 24-hour horizon optimization (K steps) using the receding horizon principle. The objective is to minimize total grid import:

$$\min_{\mathbf{P}^{batt}} \sum_{\tau=t}^{t+K} P_{grid,\tau} \tag{10}$$

subject to battery dynamics (6)–(8), power flow equations (3), and grid constraints including voltage limits ($V^{min} \le |V_{i,\tau}| \le V^{max}$) and thermal line ratings ($L_{ij,\tau} \le L^{max}$). Only the first control action $\mathbf{P}_t^{batt}$ is applied, and the horizon is re-optimized at each timestep.

TABLE I
NOMENCLATURE

Symbol	Description	Unit
$\mathcal{N}, \mathcal{E}$	Set of buses and lines	–
SOC_t	Battery state-of-charge at time t	p.u.
E_{cap}	Battery usable energy capacity	kWh
$P_{i,t}^{inj}$	Net active power injection at bus i	MW
P_t^{batt}	Battery power (positive: discharge)	kW
P^{max}	Battery maximum charging/discharging power	kW
$V_{i,t}$	Voltage magnitude at bus i	p.u.
$L_{ij,t}$	Line loading percentage for line (i,j)	%
Δt	Simulation time step (15 min)	h
H, K	Look-back and forecast horizon steps	–
$\hat{\mathbf{P}}$	Forecasted power trajectory	kW

1) Forecasting Model Architecture: The forecasting component employs Chronos 2 [31], a pre-trained Transformer-based foundation model for time-series prediction, fine-tuned on the community's historical energy-meter data using the AutoGluon-TimeSeries framework [32].

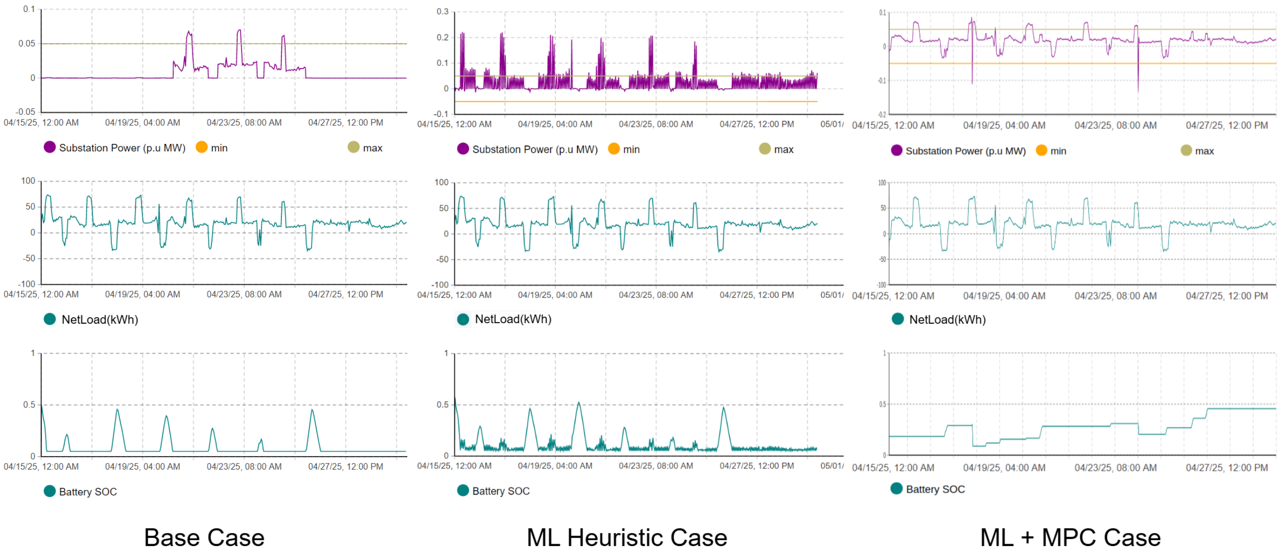

Fig. 3. Temporal comparison of the substation power, net load, and SOC of the battery for the three use cases

a) Tokenization and inference: Given a univariate time series $\mathbf{x} = \{x_1, \ldots, x_T\}$ sampled at $\Delta t = 15\,\text{min}$, Chronos 2 tokenizes each observation via mean scaling and quantile binning, then models the token sequence with a causal Transformer trained to minimise categorical cross-entropy. At inference, S trajectories are auto-regressively sampled and detokenised; the point forecast $\hat{x}_{T+h}$ is the empirical mean and quantile forecasts ($q = 0.1, \ldots, 0.9$) are the corresponding empirical quantiles.

b) Fine-tuning: The pre-trained model is fine-tuned per asset using AutoGluon's `TimeSeriesPredictor` with the Weighted Quantile Loss (WQL) as evaluation metric.

c) Temporal configuration: The lookback window spans 1 344 steps (14 days), combining the minimum lookback (672 steps) and the weekly lag distance (672 steps). The forecast horizon $H = 96$ steps covers 24 hours at 15-minute resolution, matching the battery optimisation control horizon.

d) Feature engineering: Raw meter readings are augmented with four groups of *known future covariates*: (i) temporal features (cyclical hour, day-of-week, month encodings), (ii) lag features (shifted values and rolling statistics at the 672-step weekly lag), (iii) holiday flags from Nager.Date (NL public, bank, and school holidays), and (iv) weather variables from Open-Meteo (temperature, humidity, wind, precipitation, radiation, cloud cover).

e) Alternative model families: AutoGluon-TimeSeries also supports DeepAR [33], Temporal Fusion Transformer [34], PatchTST [35], and classical statistical methods (ETS, ARIMA, Theta). Among all evaluated architectures, Chronos 2 achieved the lowest WQL across all asset time series and was therefore selected for the experiment.

2) Physics-Informed Constraint Integration: The Chronos 2 forecasting model (Section III-D1) is purely data-driven and may produce predictions that violate operational constraints such as voltage limits ($0.95 \leq |V_i| \leq 1.05$ p.u.), thermal line ratings, contracted power limits, or battery rate bounds. Three strategies were evaluated to enforce physical feasibility.

a) Strategy A — Post-processing constraint layer: Per-asset physical bounds are applied after the forecast:

$$\hat{x}_{T+h}^{\text{final}} = \text{clip}\big(\hat{x}_{T+h},\ x_i^{\min},\ x_i^{\max}\big), \tag{11}$$

guaranteeing hard constraint satisfaction without model modification.

b) Strategy B — Custom loss with physics penalty: The training loss is augmented with a penalty term:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{data}} + \lambda_{\text{phy}} \cdot \mathcal{L}_{\text{physics}}, \tag{12}$$

where $\mathcal{L}_{\text{physics}} = \sum_h \big[\text{ReLU}(\hat{x}_h - x^{\max}) + \text{ReLU}(x^{\min} - \hat{x}_h)\big]$.

c) Strategy C — Differentiable constraint injection: A differentiable power-flow solver is embedded in the training loop so the model observes grid-level consequences (voltage and line-loading violations) of its predictions end-to-end:

$$\mathcal{L}_{\text{physics}} = \sum_{h,\, i \in \mathcal{N}} \Big[\text{ReLU}\big(|V_{i,h}| - V^{\max}\big) + \text{ReLU}\big(V^{\min} - |V_{i,h}|\big)\Big] + \sum_{(i,j) \in \mathcal{E}} \text{ReLU}\big(L_{ij,h} - L^{\max}\big) \tag{13}$$

d) Experimental findings: Strategy A was straightforward to deploy with no forecast quality degradation. Strategy B degraded accuracy: the physics penalty shifted the learned token distribution away from the true data distribution (analogous to catastrophic forgetting [36]), and gradient-scale mismatch caused training instability.

IV. Simulation Setup

A. Case Study: Dutch Pilot Community

The simulation scenario represents a real-world Dutch energy community comprising 13 buses (1 slack, 12 PQ) at the LV level (400V). The network includes:

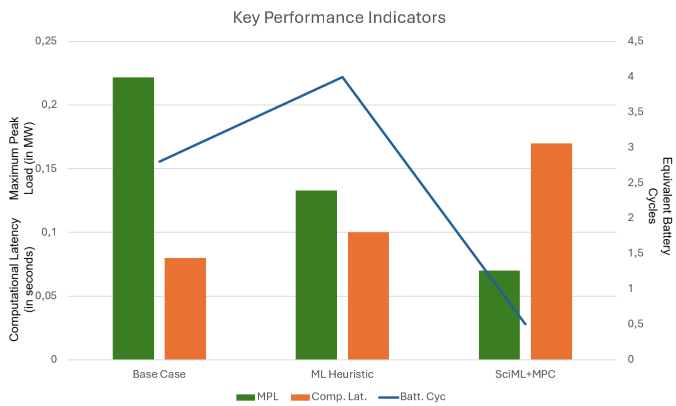

Fig. 4. Key Performance Indicators for the three use cases

- *Generation*: 5 PV installations (Real data from 3 available)
- *Storage*: 1 Centralized BESS (425 kWh, 200 kW) at the "Battery" bus
- *Loads*: Mixed residential/commercial profiles with hourly resolution
- *Transformer*: 1 MVA, 10/0.4 kV (modeled as slack bus with infinite bus approximation)
- *Cables*: Specifications provided by the industry standards and information from the demonstrator site.

B. Simulation Parameters

The discrete-event simulation runs with the following specifications:

- *Time horizon*: 24 hours (extendable to seasonal analysis)
- *Resolution*: $\Delta t = 0.25$ h (15 min), 96 timesteps/day
- *Power flow tolerance*: 10^{-5} p.u.
- *Voltage constraints*: $0.95 \leq |V_{i,t}| \leq 1.05$ p.u.
- *Maximum line loading*: $L_{ij,t} \leq 100\%$
- *Simulation time unit*: Hours (model time)

C. Key Performance Indicators

The following Four KPIs evaluate the control strategy effectiveness of three use cases:

1) Maximum Peak Load (MPL): Highest grid exchange power normalized by transformer rating (in MW):

$$\text{MPL} = \frac{\max_t |P_{\text{load}}(t) - P_{\text{gen}}(t) + P_{\text{batt}}(t)|}{S_{\text{transformer}}} \tag{14}$$

2) Computational Latency: Time from measurement to control dispatch (in seconds):

$$t_{\text{lat}} = t_{\text{pred}} + t_{\text{opt}} + t_{\text{solve}} \tag{15}$$

3) Battery Cycling: Annual equivalent full cycles (number of cycles):

$$N_{\text{cyc}} = \frac{\sum_{t=1}^{T} |P_{\text{batt}}(t)|\Delta t}{2E_{\text{rated}}} \tag{16}$$

V. Results and Discussion

A comparison of three cases have been demonstrated here:

1) Base case: Use case with only critical control on the battery.
2) ML Heuristic case: Data driven ML providing the forecast for demand and generation which in turn is fed to the heuristic rules as a control strategy for the battery serving as one of the baselines.
3) ML MPC Control case: SciML providing the forecast strategy to the MPC working as hard constraint for the SciML and determining the setpoint for the battery.

The results are demonstrated in terms of temporal series in figure 3 and KPIs in figure 4. An interesting outlook that can be seen here is that, only data driven ML forecasting with heuristic approach can be detrimental in some cases especially when there is limited real data for peak shaving as its anticipation can be sometimes not close to the reality and in turn create some oscillations. However, when the SciML with MPC approach was taken, it not only kept the peak load at the substation closer to the constraints but also demonstrated a proper DR response through the battery as shown in figure 3. The maximum peak load also improved significantly in both the cases. However, the computational latency increased with increase in sophistication of the control strategy. One interesting observation from the result was that the number of equivalent battery cycles increased in case of data driven machine learning which can be attributed to the fact that the data driven forecasting have an outlook to try to predict the future net load as it is targeted to increase the local consumption of the community. But in doing so, it sometimes adversely impact the health of the battery by increasing significantly the equivalent battery cycles.

VI. Conclusions

The present study demonstrates one of the initial real-world application of state-or-the-art Scientific Machine Learning in local energy systems and showcased how SciML can perform better than data-driven machine learning on a real demonstrator site. It also showcased that sometimes pure data driven ML algorithms can be detrimental in development of control strategies especially where there is limited data. In terms of application of SciML, Strategy C remains theoretically appealing but its implementation with Chronos 2 is hindered by the non-differentiable tokenization pipeline, which prevents gradient flow from the power-flow solver back into the model. A more pragmatic alternative is to supply grid-state variables like bus voltage magnitudes, line loadings, transformer tap positions as *known future co-variates* alongside the existing weather and calendar features. This preserves the unmodified Chronos 2 loss while giving the model physical context from which to learn operationally feasible patterns. Hard constraint satisfaction is still guaranteed by the Strategy A clipping layer downstream. Evaluating the forecast improvement provided by these physics-informed co-variates is currently being studied.

ACKNOWLEDGMENT

This work received funding from the European Union's Horizon Europe research and innovation programme under the project U2Demo, grant agreement no. 101160684 (DOI: https://doi.org/10.3030/101160684). The views and opinions expressed in this document are those of the authors only and do not necessarily reflect those of the European Union or the European Climate, Infrastructure and Environment Executive Agency (CINEA). Neither the European Union nor the granting authority can be held responsible for them.

REFERENCES

[1] Lowitzsch, J., Hoicka, C.E. and van Tulder, F.J., 2020. Renewable energy communities under the 2019 European Clean Energy Package–Governance model for the energy clusters of the future?. Renewable and Sustainable Energy Reviews, 122, p.109489.

[2] Mocci, S., Ruggeri, S. and Pilo, F., 2024. Low-voltage renewable energy communities' impact on the distribution networks. Energies, 18(1), p.126.

[3] Abu, S.M., Hannan, M.A., Rahman, S.A., Ker, P.J. and Long, C.Y., 2025. Optimization algorithms for hybrid energy storage systems based microgrid performance enhancement. Energy, 332, p.137304.

[4] L. Thurner et al., "Pandapower - an open source python tool for convenient modeling, analysis and optimization of electric power systems," *IEEE Trans. Power Syst.*, vol. 33, no. 6, pp. 6510–6521, 2018.

[5] Depeng, H., Lacarcel, F.J.S. and Simón-Moya, V., 2024. The dynamics of energy communities and innovative cooperatives: Mapping current knowledge and future trends. Journal of Innovation Knowledge, 9(4), p.100626.

[6] Ansari, A.F., Shchur, O., Küken, J., Auer, A., Han, B., Mercado, P., Rangapuram, S.S., Shen, H., Stella, L., Zhang, X. and Goswami, M., 2025. Chronos-2: From univariate to universal forecasting. arXiv preprint arXiv:2510.15821.

[7] U2Demo Horizon Europe Project, 2024-2028. U2Demo project website. Available at: https://u2demo.eu/

[8] href="https://www.flaticon.com/free-icons/battery" title="battery icons" Battery icons created by Freepik - Flaticon

[9] Open Street Map href="https://www.openstreetmap.org/copyright"

[10] href="https://www.flaticon.com/free-icons/neural-network" title="neural network icons"Neural network icons created by Freepik - Flaticon

[11] href="https://www.flaticon.com/free-icons/data-mining" title="data mining icons"Data mining icons created by Becris - Flaticon

[12] Wen, A., Wen, B., Li, J. and Xu, J., 2026. Heterogeneous Graph Neural Network with Local and Global Message Passing for AC-Optimal Power Flow Solutions. Applied System Innovation, 9(1), p.18.

[13] Hernandez-Matheus, A., Löschenbrand, M., Berg, K., Fuchs, I., Aragüés-Peñalba, M., Bullich-Massague, E. and Sumper, A., 2022. A systematic review of machine learning techniques related to local energy communities. Renewable and Sustainable Energy Reviews, 170, p.112651.

[14] Elomari, Y., Mateu, C., Marín-Genescà, M. and Boer, D., 2024. A data-driven framework for designing a renewable energy community based on the integration of machine learning model with life cycle assessment and life cycle cost parameters. Applied Energy, 358, p.122619.

[15] Pastor-Vargas, R., Rozas, W. and Manero, J., 2026. Emerging trends: Artificial intelligence, deep learning, big data, and the role of digital twins applied to energy communities. Energy Communities, pp.299-325.

[16] Liu, J., Zhou, Y., Yang, H. and Wu, H., 2022. Uncertainty energy planning of net-zero energy communities with peer-to-peer energy trading and green vehicle storage considering climate changes by 2050 with machine learning methods. Applied Energy, 321, p.119394.

[17] Karniadakis, G.E., Kevrekidis, I.G., Lu, L., Perdikaris, P., Wang, S. and Yang, L., 2021. Physics-informed machine learning. Nature Reviews Physics, 3(6), pp.422-440.

[18] Kazmi, S.N.A., Yang, F. and Sanjari, M.J., 2025, July. Enhancing Photovoltaic Power Forecasting with Physics-Informed Neural Networks for Sustainable Energy Management. In 2025 IEEE Region 10 Symposium (TENSYMP) (pp. 1-6). IEEE.

[19] Michalakopoulos, V., Pelekis, S., Kormpakis, G., Karakolis, V., Mouzakitis, S. and Askounis, D., 2024, April. Data-driven building energy efficiency prediction using physics-informed neural networks. In 2024 IEEE Conference on Technologies for Sustainability (SusTech) (pp. 84-91). IEEE.

[20] Vo, V.T., Noeiaghdam, S., Sidorov, D., Dreglea, A. and Wang, L., 2025. Solving Nonlinear Energy Supply and Demand System Using Physics-Informed Neural Networks. Computation, 13(1), p.13.

[21] Zangato, T., Osmani, A. and Alizadeh, P., 2024, January. Prior-Information Enhanced Reinforcement Learning for Energy Management Systems. In CS IT Conference Proceedings (Vol. 14, No. 2). CS IT Conference Proceedings.

[22] Perera, A.T.D. and Kamalaruban, P., 2021. Applications of reinforcement learning in energy systems. Renewable and Sustainable Energy Reviews, 137, p.110618.

[23] Qiu, L., 2025. Reinforcement Learning Approaches for Intelligent Control of Smart Building Energy Systems with Real-Time Adaptation to Occupant Behavior and Weather Conditions. Journal of Computing and Electronic Information Management, 18(2), pp.32-37.

[24] Palma, G., Guiducci, L., Stentati, M., Rizzo, A. and Paoletti, S., 2024. Reinforcement learning for energy community management: A european-scale study. Energies, 17(5), p.1249.

[25] Bose, S., Kremers, E., Mengelkamp, E.M., Eberbach, J. and Weinhardt, C., 2021. Reinforcement learning in local energy markets. Energy Informatics, 4(1), p.7.

[26] Khaskheli, S. and Anvari-Moghaddam, A., 2024. Energy trading in local energy markets: A comprehensive review of models, solution strategies, and machine learning approaches. Applied Sciences, 14(24), p.11510.

[27] Portu, J.C., Delle Femine, C., Muro, K.S., Quartulli, M. and Restelli, M., 2025, June. Limitations of Physics-Informed Neural Networks: a Study on Smart Grid Surrogation. In 2025 IEEE Kiel PowerTech (pp. 1-6). IEEE.

[28] Shan, R., Ning, H., Xu, Q., Su, X., Guo, M. and Jia, X., 2025. Physics-informed and explainable graph neural networks for generalizable urban building energy modeling. Applied Sciences, 15(16), p.8854.

[29] Moloi, K. and Addo, K., 2025. Physics-Informed Neural Networks (PINNs) for Real-Time Grid-Forming Inverter Control: Embedding Physical Laws into Deep Learning Models.

[30] Kazmi, S.N.A., Yang, F. and Sanjari, M.J., 2025, July. Enhancing Photovoltaic Power Forecasting with Physics-Informed Neural Networks for Sustainable Energy Management. In 2025 IEEE Region 10 Symposium (TENSYMP) (pp. 1-6). IEEE.

[31] A. Fatir Ansari *et al.*, "Chronos: Learning the language of time series," *arXiv preprint arXiv:2403.07815*, 2024.

[32] N. Erickson *et al.*, "AutoGluon-TimeSeries: AutoML for probabilistic time series forecasting," *arXiv preprint arXiv:2308.05566*, 2023.

[33] D. Salinas, V. Flunkert, J. Gasthaus, and T. Januschowski, "DeepAR: Probabilistic forecasting with autoregressive recurrent networks," *Int. J. Forecasting*, vol. 36, no. 3, pp. 1181–1191, 2020.

[34] B. Lim, S. O. Arik, N. Loeff, and T. Pfister, "Temporal fusion transformers for interpretable multi-horizon time series forecasting," *Int. J. Forecasting*, vol. 37, no. 4, pp. 1748–1764, 2021.

[35] Y. Nie, N. H. Nguyen, P. Sinthong, and J. Kalagnanam, "A time series is worth 64 words: Long-term forecasting with transformers," in *Proc. ICLR*, 2023.

[36] J. Kirkpatrick *et al.*, "Overcoming catastrophic forgetting in neural networks," *Proc. Nat. Acad. Sci.*, vol. 114, no. 13, pp. 3521–3526, 2017.