Economic Model Predictive Control for Household-Level Demand Response in Citizen Energy Communities

Pau Comas[a], Sebastian Madrigal[a], Antoni Morell[a], Adela Bara[b], Simona Oprea[b], Ramon Vilanova[a] and Jose Lopez Vicario[a]

[a] *IDEAS Research Group, Universitat Autònoma de Barcelona, Bellaterra, Spain, e-mail: {pau.comas@uab.cat (CA), sebastian.madrigal@uab.cat, antoni.morell@uab.cat, ramon.vilanova@uab.cat, jose.vicario@uab.cat}*

[b] *Bucharest University of Economic Studies, Department of Economic Informatics and Cybernetics, Bucharest, Romania, e-mail: {bara.adela@ie.ase.ro, simona.oprea@csie.ase.ro}*

Abstract:
Citizen Energy Communities (CECs) are entities devised for the distributed generation and sharing of renewable energy between peers, fostering cost efficiency, environmental sustainability and local ownership. CECs are rapidly growing across Europe, supported by the increasing interest in photovoltaic energy sources and the evolution (towards flexibility) of regulatory frameworks, which enable new opportunities for energy sharing among multiple members. However, for these communities to thrive, users require accessible tools in order to manage the energy flows efficiently. Despite the growing interest, current state-of-the-art energy management solutions suffer from a gap in scalability, openness, modularity and generalisability. Whilst often succeeding in conveying the feasibility and advantages of CECs' local distribution of shared renewable energy, they fail to provide flexible solutions that adapt to the diverse needs of individual users. The OPEN4CEC project, part of the Driving Urban Transitions (DUT) partnership, aims to address this need by developing an open and collaborative ecosystem of digital solutions (microservices) designed to empower community members. Under this umbrella, this paper introduces a microservice tailored for CEC users for Demand Response Management (DRM) based on an Economic Model Predictive Control (EMPC) formulation. Effective DRM at the individual household level is key for overall community efficiency, as optimising singular load profiles prevents general energy peaks and maximises the collective self-consumption rate of the CEC. The proposed framework focuses on working with limited data. We utilise a grey-box thermal model to combine physical interpretability with data-driven flexibility, relying only on historical temperature readings, available building features, and easily accessible appliance specifications. By facilitating input requirements and simplifying constraint assumptions, the framework opts for generalisability, making it more adaptable to varying household typologies and climates. The EMPC formulation generates personalised consumption recommendations for household appliances, focusing on the trade-off between economic costs, user thermal comfort, and user lifestyle.

Keywords:
Citizen Energy Communities, Demand Response Management, Economic Model Predictive Control, Grey-box Thermal Modelling, Energy Optimisation.

1. Introduction

There is an ongoing transition from the traditional relationship between utility grids and end-users to more decentralised energy systems. That brought the emergence of Citizen Energy Communities (CECs), which represent cooperative entities where citizens, small businesses, and local authorities share Distributed Energy Resources (DERs) to produce, consume, and manage renewable energy collectively. CECs empower participants to shift from being passive consumers to active prosumers, establishing cooperative microgrids that reduce reliance on traditional infrastructure and potentiate environmental sustainability. Users can share their own DERs to neighbours, or act as traditional consumers, depending on whether energy flows are localised or not, as CECs encompass many possible shapes.

For this idea to work effectively, there is a need for open tools to be provided so that newcomers, often inexperienced, can smoothly join CECs. The primary objective of these communities is to use the generated renewable energy as much as possible, lowering overall costs and the need for power from the traditional grid. Therefore, Demand Response Management (DRM) is a critical necessity. DRM aims to provide tools for users to match household energy consumption with the fluctuating availability of shared renewable generation. By monitoring controllable loads such as Heating, Ventilation, and Air Conditioning (HVAC) systems and shiftable household appliances (washing machine, dishwasher, etc), DRM maximises collective self-consumption, reduces peak demand, and minimises monetary costs for the community members.

There are many control strategies directed to automatically control those appliances or give consumption recommendations, and Model Predictive Control (MPC) has emerged as a leader. To that extent, Economic Model Predictive Control (EMPC) offers the capability to optimise energy flows dynamically by integrating weather forecasts, electricity prices, and physical building constraints while maintaining an acceptable computational burden. Nonetheless, despite proven success in residential settings, current state-of-the-art solutions face a significant gap regarding scalability, modularity, and generalisability. Many frameworks often rely on complex ”white-box” physical models that require extensive site-specific engineering or data-heavy ”black-box” machine learning (ML) models that are inaccessible to typical users [1]. On the other side, many solutions are not devised for the broader, interconnected context of an energy community, and are too specific for laboratory settings where data is easy to get, and the domain knowledge of the users is high.

This paper, developed following the goals of the OPEN4CEC project [2] (Service-oriented Open Platform for Citizen Energy Communities), introduces an open-source framework tailored for household-level DRM in CECs. The proposed system bridges the gap between sophisticated control theory and real-world applicability by combining EMPC with grey-box thermal modelling and accessible data demands. This approach shows good potential while relying only on limited historical data and basic building and appliance descriptors to dynamically parametrise the household dynamics, ensuring the option for broad adaptability across varying typologies and climates.

The specific contributions of this work are:

- **Open-Source Framework:** We deliver an accessible, open-source service [3] tailored for CEC users with individual households. It eliminates the barrier of extensive data collection by requiring only one month of historical data and standard appliance specifications, without the need for complex building descriptors or sensors.
- **Modular and Scalable Optimisation:** The solution offers highly parametrisable appliance constraints and utilises a grey-box thermal model to balance physical accuracy with broad generalisability. Additionally, its flexible objective function allows the system to shift between prioritising individual cost savings, thermal comfort, and community-level peak shaving.
- **Data-Driven Estimation of Disturbances:** To reduce system complexity and the need for intrusive sensing, the framework does not require explicit occupancy schedules or data. Instead, it integrates ML-based (XGBoost) prediction of internal heat gains as part of a lumped disturbance from available historical and dynamic data, which improves the accuracy and reliability of the EMPC optimisation.

The remainder of this paper is structured as follows: Section 2 reviews the state of the art. Section 3 presents the methodology, where we discuss the context and features of the framework. Section 4 outlines the economic MPC formulation. Section 4 describes the experimental setup and results of the proposed system applied to a case study, and Section 5 provides the conclusions.

2. State of the Art

The developments on CECs have accelerated the need for modern, dynamic control strategies for household energy management. Some approaches rely on heuristic or rule-based systems to manage devices, particularly HVAC systems and heat pumps [4]. These simpler strategies offer good generalisability but fail to dynamically adapt to weather forecasts, variable tariffs, and user behaviour. The authors have established MPC as the standard [1], and show that it significantly outperforms traditional rule-based controllers in unlocking the energy flexibility of building systems [1, 5, 6], and some have incorporated ML to capture complex building dynamics [7]. Beyond strictly economic objectives, MPC frameworks allow for combinations to minimise other important indicators such as carbon footprint [8].

A critical component in the MPC's for households is the identification of the building thermal model. High-fidelity white-box models require exhaustive physical descriptions of the building, whereas black-box models (e.g., ANNs) depend on massive historical datasets that new users typically do not have [1]. To balance model complexity and performance, grey-box models, which combine simplified physical structures such as RC networks with data-driven parameter estimation, have emerged as the ideal sweet spot.

As discussed, in EMPC, the primary objective is usually to minimise monetary costs while maintaining user comfort constraints. Early literature heavily emphasised the importance of integrating stochastic occupancy prediction to address the gap between simulated and actual building energy efficiency [9–11]. More recently, researchers have attempted to embed specific user preferences and multi-zone controls directly into the optimisation cost function [12,13]. Nonetheless, user privacy and accessibility is a concern for entities like CECs. Non-intrusive ML approaches have been successfully employed to forecast occupancy purely from power consumption data, allowing home energy systems to dynamically adjust temperature setpoints [14]. Due to computational complexity, most studies still simplify the problem by assuming fixed, rigid occupancy schedules rather than implementing active prediction [15].

Despite these theoretical advancements, a substantial gap remains regarding real-world generalisability and scalable deployment. Standard MPC solutions frequently struggle with high computational demands on legacy hardware [7] or remain confined to specific, heavily engineered simulation environments [15, 16], though recent field tests on commercial buildings with PV and batteries prove that practical EMPC deployment is viable [17]. In addition, the technical implementation in the case of CECs often relies on overly complex multi-agent systems, usually designed for optimal peer-to-peer sharing [18]. While legal frameworks across Europe are progressively empowering prosumers [19–21], these multi-agent systems struggle to onboard new buildings without excessive site-specific engineering. Additionally, the digital barrier is rarely addressed: studies show that less tech-savvy citizens are often left behind due to the complexity of these energy solutions [22]. Therefore, there is a critical need for a solution that allows CEC users to understand DRM with EMPC with a low-barrier, open-source architecture.

3. Methodology

This section details the proposed framework for DRM in households within CECs. The methodology is divided into four main components: the problem formulation within the CEC, the building thermal modelling and state estimations, the handling of forecasts and external disturbances, and finally the formulation of the EMPC optimisation.

3.1. System Architecture and Problem Formulation

To establish the operational view of this methodology, the proposed EMPC framework is first contextualised within a broader CEC. As illustrated in Figure 1, the architecture involves multiple residential prosumers interacting with local and shared energy resources. While this specific figure highlights centralised Distributed Energy Resources (DERs) at the community level, CECs can incorporate decentralised, peer-to-peer shared resources as well.

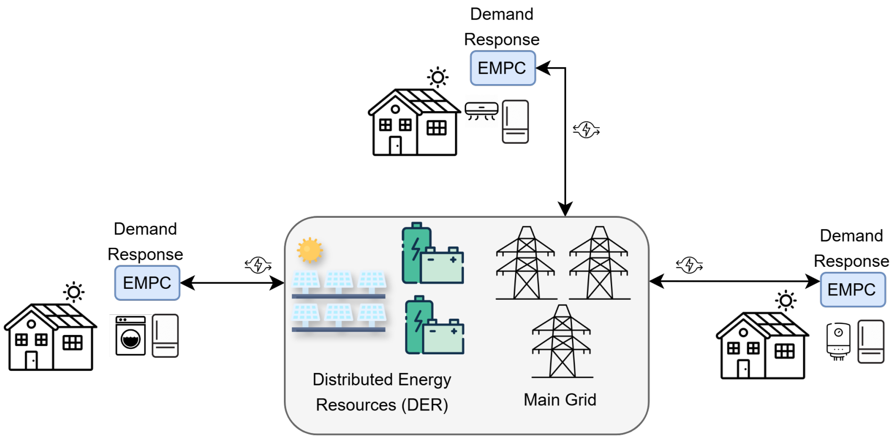

Figure 1: Architecture of the Citizen Energy Community (CEC) featuring local Economic Model Predictive Control (EMPC) and central Distributed Energy Resources (DER).

At the household level, each building is equipped with a local EMPC system. This controller serves as an edge-level DRM, interfacing directly with on-site assets such as HVAC systems and shiftable appliances (e.g., washing machines). Meanwhile, the central DER node or community aggregator coordinates shared assets such as community PV arrays and battery energy storage systems (BESS). The primary goal of the proposed local EMPC is to optimise individual household energy costs and maintain user comfort. However, its modular formulation ensures it is highly adaptable for higher-level control implementations. By distributing the computational burden to the edge, this hierarchical structure preserves user privacy regarding occupancy and device usage, while allowing objective constraints related to community-level objectives such as peak shaving and collective energy sharing.

3.2. Building Thermal Modelling and State Estimation

To optimise user consumption without compromising thermal comfort, the thermal dynamics of the building must be modelled well. A low-order grey-box resistance-capacitance (RC) network is employed to balance computational efficiency with physical accuracy. More specifically, we adopt the model structure proposed by Moritz et al. [23], which utilizes a 4R3C formulation, where 3R2C is for the building envelope. As illustrated in Figure 2, this is a system with three total thermal capacitances: two capacitances representing the inner ($C_{w,in}$) and outer ($C_{w,out}$) layers of the building envelope's thermal mass, and a third representing the indoor air mass (C_{air}). Adding further capacitance states to the wall introduces impactful complexity and risks overfitting when training data is scarce.

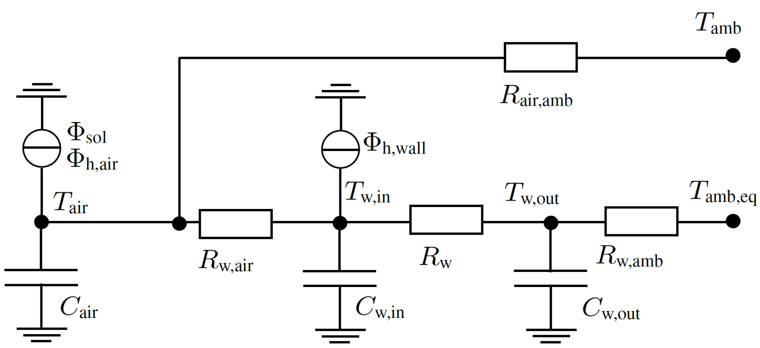

Figure 2: Schematic representation of the thermal building model featuring a 3R2C wall structure, detailing the heat fluxes (Φ), temperatures (T), resistances (R), and capacitances (C).

The differential equations of this RC circuit are governed by the physical heat transfer between the external ambient temperature (T_{amb}), the building's thermal mass ($T_{w,in}$, $T_{w,out}$), and the indoor air

(T_{air}). The external disturbances driving the system include solar radiation (Φ_{sol}) and the ambient temperature, which is split: standard T_{amb} affects the indoor air through windows and ventilation, and an equivalent ambient temperature ($T_{amb,eq}$), which combines T_{amb} and solar radiation, affects the outer wall. Also the input, represented by the HVAC, is separated into a convective part heating the air ($\Phi_{h,air}$) and a radiant part heating the inner wall ($\Phi_{h,wall}$).
The full derivation of these continuous-time differential equations and parameter definitions can be found in [23]. For the practical implementation this continuous-time model is transformed into a discrete-time linear state-space system:

$$x_{k+1} = A_d x_k + B_d u_{all,k} \quad (1)$$

$$y_k = C_d x_k + v_k \quad (2)$$

where $x_k = [T_{air}, T_{w,in}, T_{w,out}]^T$ represents the state vector comprising the three temperatures. The combined input vector $u_{all,k} = [P_{HVAC,k}, \Phi_{sol,k}, T_{amb,k}]^T$ contains both the controllable HVAC thermal power ($P_{HVAC,k}$) and the measurable external disturbances (solar radiation and ambient temperature). The output y_k is the measured indoor operative temperature, and v_k represents measurement noise.
One of the strengths is that the thermal RC parameters can be accurately identified using just one month: three weeks for training and one week for validation [23]. The required input features are strictly limited to variables readily available in most smart home setups: internally, the HVAC electrical consumption and the indoor operative temperature; externally (obtainable through owned or exogenous sources), the outdoor temperature and solar radiation.
While isolated testing of these low-order models typically yields temperature prediction errors in the 0.5-0.6°C range, real-world deployment often introduces higher errors and temperature offsets. These deviations are primarily caused by unmeasured internal heat gains (Q_{int}) resulting from human occupancy, lighting, and non-monitored appliances during the data gathering phase, jointly with other unmodeled disturbances (e.g., opening windows).
To mitigate these unmodeled disturbances while avoiding intrusive or extra hardware, a discrete-time Augmented Kalman Filter (AKF) is implemented. The model is augmented with a fourth state to estimate the unmeasured internal heat gains actively. The filter recursively estimates the unmeasurable wall temperatures and the internal heat gains while simultaneously filtering out sensor noise. The augmented state vector, denoted as $x_{aug} = [T_{air}, T_{w,in}, T_{w,out}, Q_{int}]^T$, is updated using the Kalman gain K based on the error between the measured and predicted indoor temperature:

$$\hat{x}_{aug,k|k} = \hat{x}_{aug,k|k-1} + K(y_k - C_{aug}\hat{x}_{aug,k|k-1}) \quad (3)$$

This provides the corrected initial state estimation $\hat{x}_{k|k}$ to the EMPC at the beginning of each optimisation step. Once estimated, the internal heat gain (Q_{int}) is extracted from the augmented state vector and passed to the optimiser not as an internal system state, but as a known external disturbance. It is then exponentially decayed and projected forward to improve the accuracy of the EMPC temperature predictions and not make it over-reliant on it, as it is a short-term representation, thus enabling reactivity to real-time disturbances and correcting for offset drifts caused by sudden occupancy heat gains. We prefer this over installing physical occupancy sensors, both of which introduce significant computational overhead, privacy concerns, and hardware complexity on the user side.

3.3. Forecasting and Disturbance Variables

The predictive nature of the EMPC requires accurate forecasts of exogenous variables over the prediction horizon. These disturbances are categorized into environmental variables and pricing signals, and internal gains or disturbances.

- **Environmental and Economic Signals:** Forecasts for outdoor air temperature and solar radiation are provided to the controller. In deployment, if local high-resolution datasets are unavailable, these forecasts are gathered dynamically via the Open-Meteo API [24]. To optimise operational costs, the controller requires day-ahead dynamic electricity pricing. For this framework, hourly spot market prices are retrieved from the European Network of Transmission System Operators for Electricity (ENTSO-E) transparency platform [25].

- **Occupancy and Internal Gains:** The presence of occupants directly affects internal heat gains (Q_{int}). While the AKF provides reactivity to unexpected heat gains of the moment, the EMPC could leverage a dynamic prediction of these disturbances across the entire future prediction horizon. To achieve this, we developed a novel eXtreme Gradient Boosting (XGBoost) forecasting model. Inspired by the robust data-driven methodology in [26] for estimating CEC partition coefficients, our XGBoost regressor is specifically formulated to predict future Q_{int} profiles. The model uses information available at the decision time with three distinct categories of input features:
  - *Calendar Features:* Hour of the day, day of the week, and binary flags for weekends and public holidays.
  - *Autoregressive Lags:* Historical values of previously estimated internal heat gains at 1-hour, 24-hour, and 168-hour (one week) intervals.
  - *Temperature Features:* Current indoor operative temperature and forecasted outdoor temperature.

  In Section 4.2. we will evaluate the impact of this data-driven forecasting strategy, comparing the XGBoost against the AKF correction and applying the exponential decay to the current Q_{int} value estimated by the filter.

3.4. Economic Model Predictive Control (EMPC) Formulation

The core of the local framework is a Mixed-Integer Quadratic Programming (MIQP) optimisation for the recommendations of activation of the HVAC system and shiftable appliances. The problem is formulated using the CVXPY modelling language and solved via the Gurobi optimizer. This combination provides the necessary computational speed and robustness within a receding horizon framework, ensuring real-time operational feasibility at the edge.

3.4.1. Objective Function

The multi-objective cost function J is evaluated over the discrete prediction horizon N. It is designed to minimize the economic operational cost of electricity, penalize grid power peaks, and ensure thermal comfort:

$$\min_{u,z,\epsilon} \sum_{k=0}^{N-1} \left[C_k^{grid} P_{total,k} \Delta t + \gamma P_{total,k}^2 + \lambda \left(\epsilon_{lower,k} + \epsilon_{upper,k} \right) \right] \tag{4}$$

where C_k^{grid} is the dynamic day-ahead electricity tariff, and $P_{total,k} = P_{HVAC,k} + \sum_{i \in \mathcal{S}} P_{shift,i,k}$ represents the total controllable power consumed by the HVAC and the set of shiftable appliances $\mathcal{S}$. To ensure economic accuracy, the true total cost of non-interruptible appliance cycles that extend beyond the immediate prediction horizon is pre-calculated and assigned to the initiation variable.

A key advantage of this formulation is its parametrisability. The quadratic penalty factor γ introduces a peak-shaving component, because tuning this parameter higher forces the optimiser to flatten the load profile, reducing sudden spikes in grid demand, which is beneficial for the wider energy community. Conversely, the linear slack penalty λ controls the rigidity of thermal comfort bounds that will be introduced in 3.4.2. Tuning λ allows the system to balance aggressive cost savings (relaxing λ) against strict occupant comfort enforcement (increasing λ).

3.4.2. System Constraints

The optimisation is subjected to the following operational and physical constraints:

1. *Thermal Dynamics and Comfort Bounds:* The future states of the building are predicted using

the discrete state-space model. The internal heat gains $Q_{int,k}$ are explicitly incorporated as a predicted disturbance affecting the indoor air temperature state:

$$x_{k+1} = A_d x_k + B_d \begin{bmatrix} P_{HVAC,k} \\ \Phi_{sol,k} \\ T_{amb,k} \end{bmatrix} + EQ_{int,k} \tag{5}$$

where E is the discrete disturbance vector derived from the continuous-time mapping $[\frac{1}{C_{air}}, 0, 0]^T$, injecting the heat gains directly into the indoor air mass dynamics. Depending on the chosen forecasting methodology, the vector $Q_{int,k}$ across the horizon is either populated by the data-driven XGBoost predictions or approximated by an exponential decay of the current real-time state estimated by the Kalman filter.

Additionally, the predicted indoor air temperature $T_{\text{air},k}$ (extracted from the state vector x_k) must remain within the dynamic, occupancy-driven lower ($T_{\text{min},k}$) and upper ($T_{\text{max},k}$) bounds:

$$T_{\text{min},k} - \epsilon_{\text{lower},k} \leq T_{\text{air},k} \leq T_{\text{max},k} + \epsilon_{\text{upper},k} \quad \text{with} \quad \epsilon_{\text{lower},k}, \epsilon_{\text{upper},k} \geq 0 \tag{6}$$

2. *HVAC Actuator and Compressor Protection:* The HVAC system operates under strict customisable capacity limits (P_{min}, P_{max}). A binary variable $z_{HVAC,k} \in \{0, 1\}$ dictates the on/off status of the unit:

$$P_{min} z_{HVAC,k} \leq P_{HVAC,k} \leq P_{max} z_{HVAC,k} \tag{7}$$

To prevent compressor damage from rapid short-cycling, a minimum off-time constraint is enforced. If the HVAC transitions from on to off, it must remain off for a designated parametrizable cooldown period t_{cool}:

$$z_{HVAC,k+j} \leq 1 - (z_{HVAC,k-1} - z_{HVAC,k}) \quad \forall j \in \{1, \dots, t_{cool}\} \tag{8}$$

3. *Generalised Shiftable Appliance Logic:* The framework supports the integration of one or multiple shiftable appliances $i \in \mathcal{S}$ (e.g., washing machines, dishwashers). These non-interruptible loads are modelled using a predefined discrete power profile W_i of length L_i. This array captures multi-phase appliance behaviours. A binary decision variable $z_{start,i,k} \in \{0, 1\}$ dictates the cycle's initiation. The appliance's total power draw at any time step is the convolution of historical start decisions and its specific power profile:

$$P_{shift,i,k} = \sum_{j=0}^{L_i - 1} z_{start,i,k-j} W_{i,j} \quad \forall i \in \mathcal{S} \tag{9}$$

To respect user preferences, each appliance is constrained to start only within a user-defined schedule slot (e.g., specific hours of the day or chosen days of the week), mapped by a customisable binary parameter array $\mathcal{M}_{valid,i,k}$. A global constraint ensures the cycle executes exactly the required number of times R_i during the scheduling horizon:

$$z_{start,i,k} \leq \mathcal{M}_{valid,i,k} \quad \forall i \in \mathcal{S} \tag{10}$$

$$\sum_{k=0}^{N-1} z_{start,i,k} = R_i \quad \forall i \in \mathcal{S} \tag{11}$$

4. Case Study: PLEGMA

4.1. Experimental Setup

The case study utilizes the PLEGMA dataset [27], a Greek open-source residential electricity dataset that includes high-resolution measurements of both aggregate and appliance-level consumption across

13 households. The dataset provides environmental data, such as indoor and outdoor temperature and humidity, alongside metadata detailing appliance specifications (e.g., wattage thresholds, minimum active/inactive times). This makes PLEGMA highly suitable for our use-case.
Electrical measurements were resampled to 15-minute averages to align with the indoor environmental measurements, and is the EMPC sampling time (T_s). The prediction horizon is 24 hours, applying control actions in a standard receding-horizon format. House 2 during the heating season is used for displaying the results, encompassing a four-week period from February 24, 2023, to March 24, 2023 with the last week being the used for validation.
To demonstrate load-shifting, the scenario schedules a washing machine twice weekly. The controller optimises only the start time, restricted to a user-preferred window between 15:00 and 23:45 on designated days. Once triggered, the deterministic 3-hour cycle unfolds exactly once per active day, progressing through three 1-hour phases: 700 W, 75 W, and 100 W. For the HVAC system, the baseline parameters derived from the dataset dictate an operational power range between a minimum of 100 W and a maximum of 1600 W. Additionally, the controller enforces a minimum OFF time of 45 minutes to prevent excessive short-cycling of the compressor.

4.2. Results

As illustrated in Figure 3, the historical data of PLEGMA's house 2 reveals a distinct user lifestyle pattern. The property is predominantly occupied between 15:00 and 24:00, during which the user actively engages the heating system. Outside of this time frame, the heating system remains largely inactive, resulting in lower indoor temperatures.

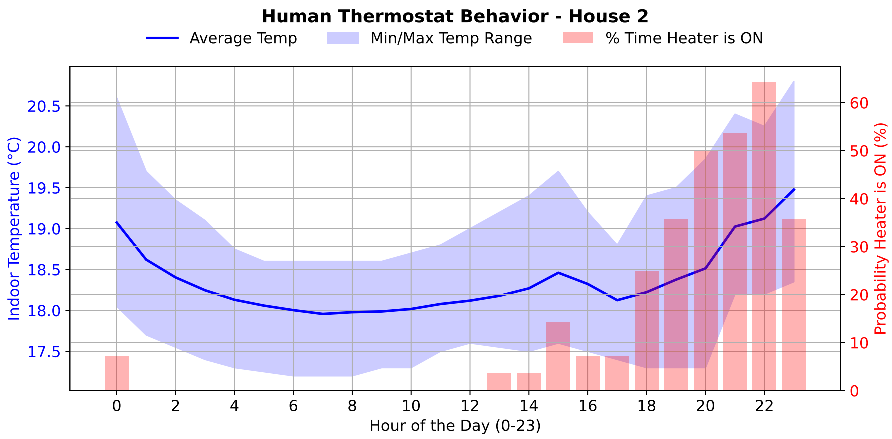

Figure 3: Human Thermostat Behavior - Historical Data Analysis.

The first operational scenario evaluates the EMPC's ability to respect and optimise this specific user lifestyle. Dynamic comfort bounds were established, featuring a tight lower bound set at $0.2°C$ below the historical user temperature and a relaxed upper bound of $24°C$. Figure 4 displays the system's results under these conditions. The EMPC successfully manages the thermal dynamics, utilising the available flexibility to pre-heat when energy is cheaper while getting closely to the user's preferred thermal profile.
Table 1 shows the aggregate performance metrics. Under what we call the "user lifestyle" bounds, the EMPC achieves both energy and economic savings. Total energy consumption is reduced from 22.48 kWh to 21.90 kWh, yielding a cost reduction of 2.5% (from €3.19 to €3.11). Furthermore, the predictive controller effectively helps reduce peak energy demand, lowering the maximum power peak by 10.4% (from 1993.1 W to 1786.1 W). The EMPC registers 13 comfort violations out of 672, a minor trade-off, which occurs as the optimiser strategically navigates the extreme edges of the tight lower bound to maximise economic savings, and this can be shifted with parametrisation.
The second scenario tests the EMPC against default thermal comfort bounds from World Health Organisation (WHO) [28]. These guidelines label minimum a value of $18°C$ during generally un-

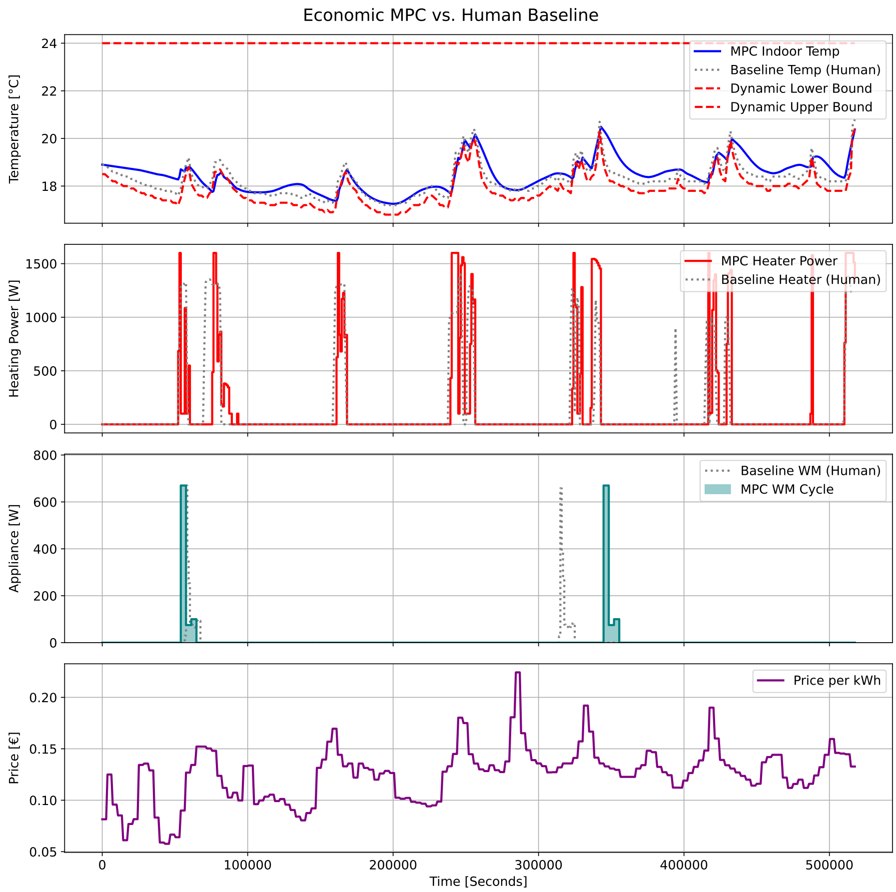

Figure 4: EMPC vs. Human Baseline: User Lifestyle Bounds.

Table 1: Operational Performance: User Lifestyle Bounds

Metric	Human Baseline	Economic MPC	Difference
Total Energy (kWh)	22.48	21.90	-2.6%
Total Cost (€)	3.19	3.11	-2.5%
Peak Energy (W)	1993.1	1786.1	-10.4%
Comfort Violations	0	13	+13

occupied or sleeping hours (00:00–15:00) on cold seasons, and $20°C$ during active occupancy hours (15:00–24:00). As demonstrated by the historical baseline (Figure 3), the manual operation of the HVAC system by the user fails to maintain these standards, resulting in a substantial 336 comfort violations (3.5 days).

When the EMPC is tasked with enforcing these WHO bounds (Figure 5), it drastically improves the thermal environment, reducing comfort violations by 84.8%, and it could enforce further comfort violations removal by being more strict on the soft constraints. As shown in Table 2, it also means an increase in heating effort. Consequently, total energy consumption rises by 65.4% (to 37.18 kWh), and total operational costs increase by 46.1% (to €4.66). Again, peak shaving is accomplished due to washing machine rescheduling, achieving a 19.7% reduction in peak grid stress compared to the human baseline's 1993.1 W.

Table 2: Operational Performance: WHO Standard Bounds ($18°C$ - $20°C$)

Metric	Human Baseline	Economic MPC	Difference
Total Energy (kWh)	22.48	37.18	+65.4%
Total Cost (€)	3.19	4.66	+46.1%
Peak Energy (W)	1993.1	1600.0	-19.7%
Comfort Violations	336	51	-84.8%

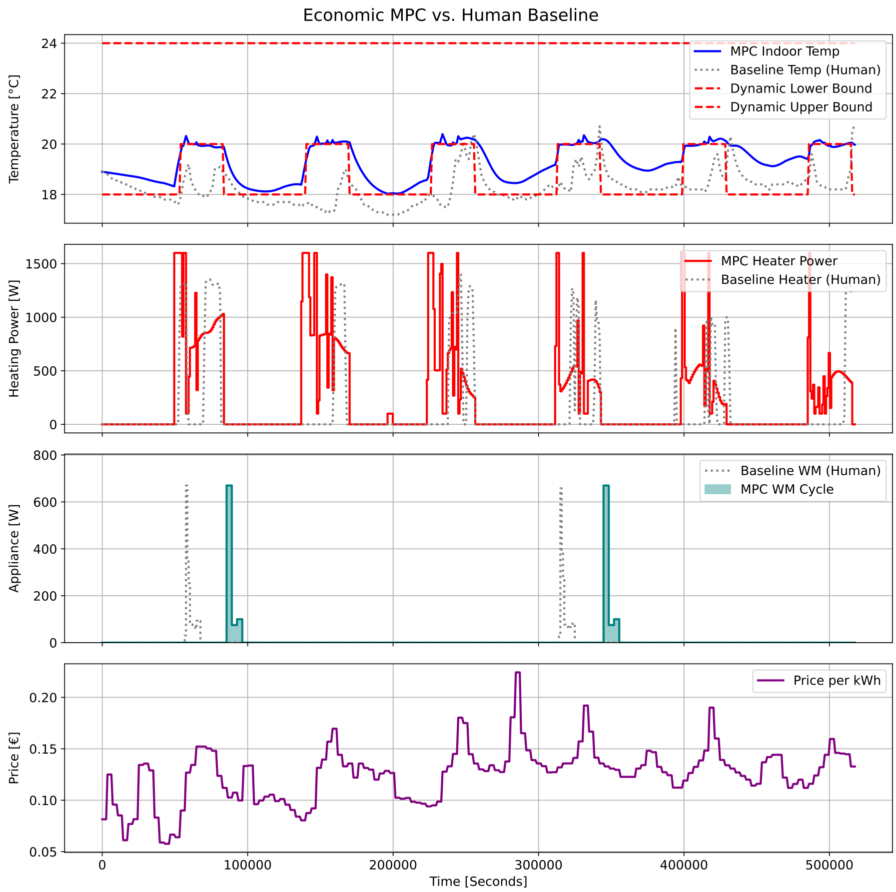

Figure 5: EMPC vs. Human Baseline: WHO Standard Bounds.

Table 3: Impact of XGBoost Heat Gain Forecasting on EMPC Performance

Scenario	Metric	With AKF	with XGBoost	Change
User Lifestyle	Total Cost (€)	3.11	2.99	-3.9%
	Comfort Violations	13	11	-15.4%
WHO Bounds	Total Cost (€)	4.66	4.63	-0.6%
($18°C$ - $20°C$)	Comfort Violations	51	32	-37.3%

Finally, to validate the impact of the data-driven forecasting strategy proposed in the methodology, we compare the model with the AKF against the XGBoost. The inclusion of the XGBoost model allows the EMPC to proactively anticipate the disturbances rather than reactively mitigating them. As seen in Table 3, this gives some improvements in both economic cost and thermal comfort. Under the first bounds, the XGBoost integration further reduces the total operational cost to €2.99 (an additional 3.9% improvement over the standard EMPC) while decreasing transient comfort violations from 13 to 11. The impact is stronger under the WHO Thermal Bounds. While the economic cost sees a marginal reduction of 0.6% (from €4.66 to €4.63), the proactive pre-heating enabled by the XGBoost forecasts drastically reduces comfort violations from 51 down to 32, a 37.3% improvement in thermal comfort compared to the AKF.

5. Conclusion

This paper presented and validated the design of an EMPC framework for household-level DRM within CECs. The key advantage of the proposed framework is that it is generalisable because of data demands and modelling complexity, and in addition, is shared with the community. The system can be accurately deployed in any household using only one month of historical indoor temperature data and standard appliance specifications. Extensive on-site sensing is avoided, as necessary variables

such as external temperature and solar radiation are adapted to be obtained via the discussed external APIs, which is a typical practice.
Furthermore, the parametrisable nature of the designed optimisation allows the system to be dynamically tailored to distinct operational goals. The framework can seamlessly shift its priority between maximising individual economic cost savings, enforcing occupant thermal comfort, or adapting peak shaving priorities due to community-level demands. Future work will focus on scaling the framework by coupling it with higher-level CEC demand response coordination. This will involve integrating the controller with local and shared energy resources, such as photovoltaic generation and battery energy storage systems, which are key elements of CEC architectures. Thanks to the modular design proposed in the code, the framework is prepared to ease the integration of these resources as new blocks, to facilitate interconnection and optimisation with cooperative microgrids.

Acknowledgments

This work has received support from the Catalan Government under Project 2021 SGR 00197 and also by the Spanish Government under MICINN projects PID2024-156522OB-C33 co-funded with the European Union ERDF funds and MCIN/AEI/10.13039/501100011033 project TED2021-129134BI00 co-funded with the European Union "NextGenerationEU"/PRTR funds, and projects OPEN4CEC PCI2023-145975-2 and SMART-LEM PCI2023-145977-2 funded by MCIU/AEI/10.1303-9/501100011033 and co-funded by European Union.

References

[1] J. Drgoňa, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, and L. Helsen, "All you need to know about model predictive control for buildings," *Annual Reviews in Control*, vol. 50, pp. 190–232, 2020.

[2] OPEN4CEC Consortium, "OPEN4CEC: Service-oriented open platform for citizen energy communities." Driving Urban Transitions (DUT) Partnership Project, 2023. Accessed: 2026-02-11.

[3] P. Comas, "Replication Assets for: "Economic Model Predictive Control for Household-Level Demand Response in Citizen Energy Communities"." CORA-RDR, 2026. DOI: 10.34810/data3143. `https://doi.org/10.34810/data3143`.

[4] T. Dengiz, P. Jochem, and W. Fichtner, "Demand response with heuristic control strategies for modulating heat pumps," *Applied Energy*, vol. 238, pp. 1346–1360, 2019.

[5] T. Q. Péan, J. Salom, and R. Costa-Castelló, "Review of control strategies for improving the energy flexibility provided by heat pump systems in buildings," *Journal of Process Control*, vol. 74, pp. 35–49, 2019.

[6] Y. Xin, Z. Zhang, Y. Zhou, Y. Liu, D. Wang, and S. Nan, "A comprehensive review of predictive control strategies in heating, ventilation, and air-conditioning (hvac): Model-free vs model," *Journal of Building Engineering*, vol. 94, p. 110013, 2024.

[7] J. Drgoňa, D. Picard, M. Kvasnica, and L. Helsen, "Approximate model predictive building control via machine learning," *Applied Energy*, vol. 218, pp. 199–216, 2018.

[8] T. Q. Péan, J. Ortiz, J. Salom, and R. Costa-Castelló, "Price and carbon-based energy flexibility of residential heating and cooling loads using model predictive control," *Sustainable Cities and Society*, vol. 50, p. 101579, 2019.

[9] A. Mirakhorli and B. Dong, "Occupancy behavior based model predictive control for building indoor climate—a critical review," *Energy and Buildings*, vol. 129, pp. 499–513, 2016.

[10] C. D. Korkas, S. Baldi, I. Michailidis, and E. B. Kosmatopoulos, "Occupancy-based demand response and thermal comfort optimization in microgrids with renewable energy sources and energy storage," *Applied Energy*, vol. 163, pp. 93–104, 2016.

[11] F. Oldewurtel, D. Sturzenegger, and M. Morari, "Importance of occupancy information for building climate control," *Applied Energy*, vol. 101, pp. 521–532, 2013.

[12] M. Frahm, P. Zwickel, J. Wachter, F. Langner, P. Strauch, J. Matthes, and V. Hagenmeyer, "Occupant-oriented economic model predictive control for demand response in buildings," in *Proceedings of the Thirteenth ACM International Conference on Future Energy Systems (e-Energy)*, pp. 354–360, ACM, 2022.

[13] M. Frahm, T. Dengiz, P. Zwickel, H. Maaß, J. Matthes, and V. Hagenmeyer, "Occupant-oriented demand response with multi-zone thermal building control," *Applied Energy*, vol. 347, p. 121454, 2023.

[14] R. Nematirad, M. M. Ardehali, A. Khorsandi, and A. Mahmoudian, "Optimization of residential demand response program cost with consideration for occupants thermal comfort and privacy," *IEEE Access*, vol. 12, pp. 15194–15206, 2024.

[15] F. Langner, W. Wang, M. Frahm, and V. Hagenmeyer, "Model predictive control of distributed energy resources in residential buildings considering forecast uncertainties," *Energy and Buildings*, vol. 303, p. 113753, 2024.

[16] P. Strauch, W. Wang, and F. Langner, "Model predictive control for demand flexibility of a residential building with multiple distributed energy resources," *Energy and Buildings*, vol. 305, p. 113753, 2024.

[17] K. Zhang, A. Prakash, L. Paul, D. Blum, P. Alstone, J. Zoellick, R. Brown, and M. Pritoni, "Model predictive control for demand flexibility: Real-world operation of a commercial building with photovoltaic and battery systems," *Advances in Applied Energy*, vol. 7, p. 100099, 2022.

[18] C. Feng and A. L. Liu, "Peer-to-peer energy trading of solar and energy storage: A networked multiagent reinforcement learning approach," *arXiv preprint arXiv:2401.13947*, 2024.

[19] S. Ahmed, A. Ali, and A. D'angola, "A review of renewable energy communities: concepts, scope, progress, challenges, and recommendations," *Sustainability*, vol. 16, no. 5, p. 1749, 2024.

[20] E. Cocciolo, "The role of energy communities for thermal networks: An eu legal perspective," *Review of European, Comparative & International Environmental Law*, vol. 33, no. 3, pp. 494–506, 2024.

[21] G. Yiasoumas, L. Berbakov, V. Janev, A. Asmundo, E. Olabarrieta, A. Vinci, G. Baglietto, and G. E. Georghiou, "Key aspects and challenges in the implementation of energy communities," *Energies*, vol. 16, no. 12, p. 4703, 2023.

[22] S.-V. Oprea, A. Bara, B. G. Tudorica, M. I. Calinoiu, and M. A. Botezatu, "Insights into demand-side management with big data analytics in electricity consumers behaviour," *Computers and Electrical Engineering*, vol. 89, p. 106902, 2021.

[23] M. Frahm, F. Langner, P. Zwickel, J. Matthes, R. Mikut, and V. Hagenmeyer, "How to derive and implement a minimalistic rc model from thermodynamics for the control of thermal parameters for assuring thermal comfort in buildings," in *Open Source Modelling and Simulation of Energy Systems (OSMSES)*, pp. 1–6, IEEE, 2022.

[24] P. Zippenfenig, "Open-Meteo.com Weather API." Zenodo, 2023. `https://doi.org/10.5281/zenodo.7970649`.

[25] EnergieID, "entsoe-py: Python client for the entso-e api." `https://github.com/EnergieID/entsoe-py`, 2024.

[26] S. Madrigal, R. Gallinad, J. L. Vicario, A. Morell, and R. Vilanova, "Improving energy distribution in collective self-consumption via XGBoost-based allocation coefficients prediction," *Applied Energy*, vol. 409, p. 127469, 2026.

[27] S. Athanasoulias, M. Bimpas, R. Besseau, S. Delbecq, F. Amato, N. Gkonis, L. Stankovic, and V. Stankovic, "The plegma dataset: domestic appliance-level and aggregate electricity demand with metadata from greece," *Scientific Data*, vol. 11, no. 1, p. 370, 2024.

[28] World Health Organization. Regional Office for Europe, "Health impact of low indoor temperatures: report on a who (world health organization) meeting held at copenhagen on november 11-14, 1985," Tech. Rep. EH-16, PB-88-163753/XAB, World Health Organization, Copenhagen, Denmark, 1987. OSTI ID: 5282879.