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 18 May 2026

Self-learning adaptive heating curve adjustment for intuitive optimization

AHA

Source: own illustration

Date: 18.05.2026

Location: Dübendorf/Grabs

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

Subsidy Recipients:
Eidgenössische Material Prüfungsanstalt (EMPA)
Urban Energy System LAB (UES LAB)
Ueberlandstrasse 129, 8600 Dübendorf
www.empa.ch

Autors:
Michael Locher, Urban Energy System Lab @ EMPA, michael.locher@empa.ch

SFOE project coordinators:
Andreas Eckmanns, andreas.eckmanns@bfe.admin.ch
Martin Ménard, menard@lowtechlab.ch

SFOE contract number: SI/502672-01

The authors bear the entire responsibility for the content of this report and for the conclusions drawn therefrom.

Summary

This research focuses on optimizing heating systems in buildings to reduce energy consumption and CO2 emissions. In Switzerland, buildings account for a large portion of the nation's energy use, with heating representing the dominant share. Traditional heating systems rely on static heating curves, configured at installation, which often become outdated over time and lead to systematic inefficiencies. This project introduces an adaptive approach that continuously adjusts the heating curve to optimize energy consumption while maintaining thermal comfort.

The main objective is to autonomously adjust the heating curve based on minimal data inputs — room temperature, outdoor temperature, and heating energy — without requiring an explicit building model. The system minimizes a score function combining normalized energy consumption and comfort deviations. Gaussian Process regression surrogate models are employed to find, via Contextual Bayesian Optimization, the optimal heating curve parameters over time, with outdoor temperature serving as the context variable that enables the algorithm to learn a context-dependent policy rather than a single fixed setpoint.

Simulations of the heating system, based on detailed EnergyPlus models for Bülsweg — a multi-family residential building — and NEST Sprint — a multi-office research unit — were conducted to validate the approach and calibrate hyperparameters prior to field deployment. These simulations use real building data, including measurements from heat meters and sensors, to assess heating energy usage and comfort levels under controlled conditions.

Field tests were conducted at two buildings with distinct typologies: Bülsweg (multi-family residential, in collaboration with Lippuner AG, Buchs) and NEST Sprint (office unit, Empa Dübendorf). At Bülsweg, the adaptive approach achieved energy savings of approximately 4–6% compared to the static baseline while maintaining nearly unchanged comfort levels (Comfort Score: 0.24 → 0.23). Among the four evaluation methods applied — Linear Regression, Heating Degree Day normalization, Machine Learning prediction, and Digital Twin analysis — the HDD-based estimate proved most reliable, balancing weather normalization and model robustness. Although the Digital Twin failed to reproduce these gains due to limited extrapolation beyond its training range (30–45°C), real-world data confirmed measurable efficiency improvements. The Bayesian Optimization model demonstrated stable convergence after 30–40 iterations, but robustness tests revealed asymmetric sensitivity: tolerant to strong positive outliers yet vulnerable to negative ones, highlighting the need for improved noise handling in future iterations.

At NEST Sprint, the evaluation required a dedicated seasonal correction: the static reference period (October–November 2025) and the adaptive phase (February–March 2026) fall in different seasons, with February–March exhibiting approximately double the solar irradiance of October–November. Two historical control periods — February–March 2024 and 2025, both verified as pure static operation — were used to isolate the seasonal baseline shift of 10–12% from the algorithm effect. After removing this confounder, the net AHA-attributable saving converges to 5-10% across two independent methods, with a central value of approximately 7.5%. Thermal comfort was maintained throughout, and comfort variability reduced by 86% (maximum comfort violation: 2.0 → 0.38). The algorithm demonstrated statistically significant context-adaptive behavior confirming that the GP surrogate learned a physically correct heating policy. A reward computation bug active during the deployment was identified post-hoc; counterfactual analysis demonstrated only marginal impact on action selection in the operating temperature range, due to the context-dependent structure of the Bayesian Optimization.

Future work will focus on scaling and refining the adaptive heating framework to enable autonomous, data-driven optimization across multiple buildings. The next development phase aims to automate data processing and model retraining, enhance robustness against sensor noise and outliers, and expand the optimization range to fully capture the system's energy-saving potential. A full-season deployment at both buildings is planned for Winter 2026/27, with the Yc hardware coupling at Bülsweg resolved and an automated sensor health monitoring system in place. Additional field tests, including applications in single-family houses with heat pumps, will assess the method's transferability and reliability under diverse operating conditions. Furthermore, integrating additional contextual factors — such as wind, solar

radiation, and occupancy dynamics — directly into the learning model will improve both predictive accuracy and comfort control. These developments will strengthen the adaptive heating approach as a scalable and robust alternative to static control strategies, advancing the transition toward intelligent, energy-efficient building operation.

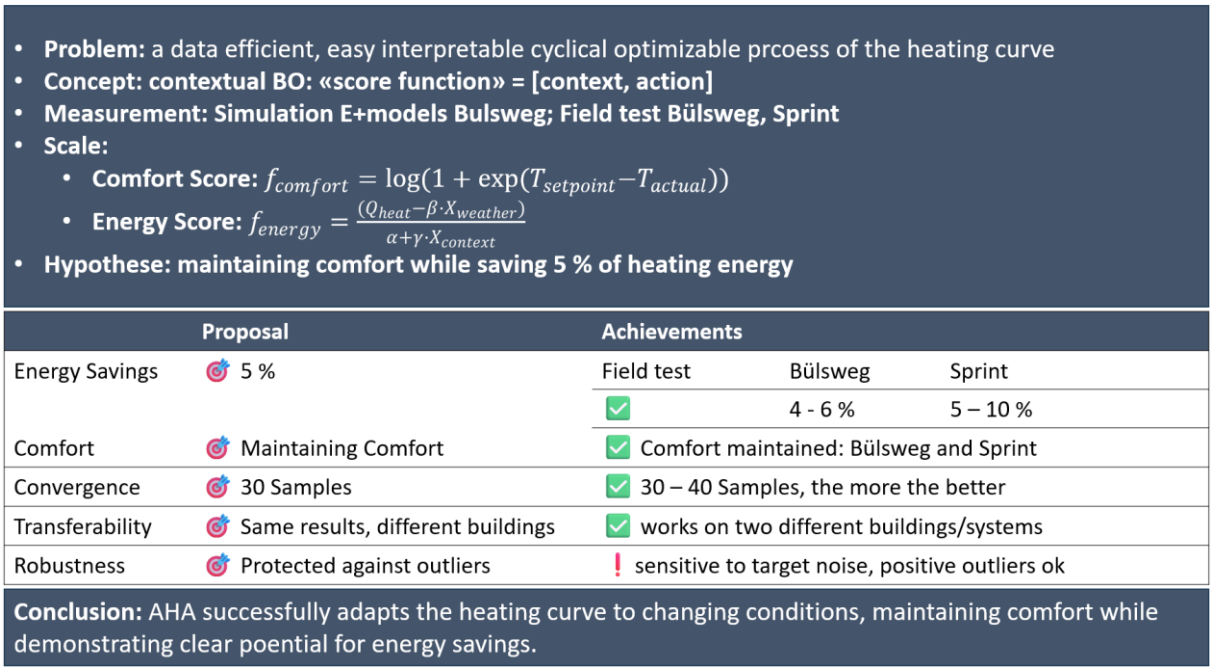

Figure 1: This figure provides a comprehensive overview of Project AHA, illustrating the core concept of optimizing heating curves using contextual Bayesian optimization. It visualizes the problem statement, the proposed approach, and the underlying hypothesis. The diagram also outlines the measurement strategy, the scale of the study, and both the proposed and achieved results. The optimization process is based on a Gaussian Process Regression model that is iteratively updated with contextual data (outdoor temperature) and system actions (inlet temperature) to minimize a score function. This score combines normalized values of energy consumption and comfort deviations relative to a chosen reference point, demonstrating how the adaptive method balances comfort maintenance with energy savings in practice.

Zusammenfassung

Diese Forschung befasst sich mit der Optimierung von Heizsystemen in Gebäuden zur Reduktion des Energieverbrauchs und der CO_2-Emissionen. In der Schweiz entfällt ein erheblicher Teil des nationalen Energieverbrauchs auf Gebäude, wobei die Wärmeerzeugung den dominanten Anteil ausmacht. Traditionelle Heizsysteme basieren auf statischen Heizkurven, die bei der Installation konfiguriert werden und mit der Zeit häufig veralten, was zu systematischen Ineffizienzen führt. Das vorliegende Projekt stellt einen adaptiven Ansatz vor, der die Heizkurve kontinuierlich anpasst, um den Energieverbrauch zu optimieren und gleichzeitig den thermischen Komfort zu erhalten.

Das Hauptziel besteht darin, die Heizkurve autonom anzupassen, basierend auf einer minimalen Datengrundlage — Raumtemperatur, Aussentemperatur und Heizenergie — ohne dass ein explizites Gebäudemodell erforderlich ist. Das System minimiert eine Bewertungsfunktion, die den normalisierten Energieverbrauch und Komfortabweichungen kombiniert. Über Contextual Bayesian Optimization werden mit Hilfe von Gaussian-Process-Regressionsmodellen die optimalen Heizkurvenparameter über die

Zeit gefunden. Die Aussentemperatur dient dabei als Kontextvariable, die es dem Algorithmus ermöglicht, eine kontextabhängige Strategie zu erlernen, statt einen einzelnen festen Sollwert zu verwenden.

Vor dem Feldeinsatz wurden Simulationen des Heizsystems durchgeführt, basierend auf detaillierten EnergyPlus-Modellen für Bülsweg — ein Mehrfamilienhaus — und NEST Sprint — eine Mehrbüro-Forschungseinheit. Diese Simulationen nutzen reale Gebäudedaten, einschliesslich Messwerten von Wärmezählern und Sensoren, um den Heizenergieverbrauch und das Komfortniveau unter kontrollierten Bedingungen zu beurteilen, und dienen der Validierung des Ansatzes sowie der Kalibrierung der Hyperparameter.

Feldtests wurden in zwei Gebäuden mit unterschiedlicher Typologie durchgeführt: Bülsweg (Mehrfamilienhaus, in Zusammenarbeit mit der Lippuner AG, Buchs) und NEST Sprint (Büroeinheit, Empa Dübendorf). In Bülsweg erzielte der adaptive Ansatz Energieeinsparungen von rund 4–6 % gegenüber der statischen Referenz, bei nahezu unverändertem Komfortniveau (Comfort Score: 0.24 → 0.23). Von den vier angewandten Auswertungsmethoden — lineare Regression, HDD-Normalisierung (Heating Degree Days), Machine-Learning-Prognose und Digital-Twin-Analyse — erwies sich die HDD-basierte Schätzung als am verlässlichsten, da sie eine ausgewogene Kombination aus Witterungsnormalisierung und Modellrobustheit bietet. Während der Digital Twin diese Einsparungen aufgrund einer begrenzten Extrapolation jenseits seines Trainingsbereichs (30–45 °C) nicht reproduzieren konnte, bestätigten die realen Messdaten die messbaren Effizienzgewinne. Das Bayesian-Optimization-Modell zeigte eine stabile Konvergenz nach 30–40 Iterationen; Robustheitstests offenbarten jedoch eine asymmetrische Empfindlichkeit: tolerant gegenüber starken positiven Ausreissern, aber anfällig für negative — ein Hinweis auf den Bedarf eines verbesserten Rauschhandlings in zukünftigen Iterationen.

Bei NEST Sprint erforderte die Auswertung eine dedizierte saisonale Korrektur: Die statische Referenzphase (Oktober–November 2025) und die adaptive Phase (Februar–März 2026) fallen in unterschiedliche Jahreszeiten, wobei die Solarstrahlung im Februar–März etwa doppelt so hoch liegt wie im Oktober–November. Zwei historische Kontrollperioden — Februar–März 2024 und 2025, beide als reiner statischer Betrieb verifiziert — wurden herangezogen, um die saisonale Verschiebung der Baseline von 10–12 % vom Algorithmuseffekt zu trennen. Nach Bereinigung dieses Confounders konvergiert die AHA-zurechenbare Netto-Einsparung auf 5–10 % über zwei unabhängige Methoden hinweg, mit einem Zentralwert von etwa 7.5 %. Der thermische Komfort wurde durchgehend aufrechterhalten, wobei die Komfortvariabilität um 86 % reduziert wurde (maximale Komfortverletzung: 2.0 → 0.38). Der Algorithmus zeigte statistisch signifikantes kontextadaptives Verhalten, was bestätigt, dass das GP-Surrogat eine physikalisch korrekte Heizstrategie erlernt hat. Ein während des Deployments aktiver Bug in der Reward-Berechnung wurde nachträglich identifiziert; eine kontrafaktische Analyse zeigte aufgrund der kontextabhängigen Struktur der Bayesian Optimization nur marginale Auswirkungen auf die Aktionswahl im operierenden Temperaturbereich.

Die zukünftige Arbeit wird sich auf die Skalierung und Verfeinerung des adaptiven Heizframeworks konzentrieren, um eine autonome, datengetriebene Optimierung in mehreren Gebäuden zu ermöglichen. Die nächste Entwicklungsphase zielt darauf ab, die Datenverarbeitung und das Modell-Retraining zu automatisieren, die Robustheit gegenüber Sensorrauschen und Ausreissern zu erhöhen und den Optimierungsbereich zu erweitern, um das volle Energieeinsparpotenzial des Systems auszuschöpfen. Ein vollständiger Saison-Einsatz an beiden Gebäuden ist für den Winter 2026/27 geplant, mit gelöster Yc-Hardware-Kopplung in Bülsweg und einem automatisierten Sensor-Health-Monitoring. Zusätzliche Feldtests, einschliesslich Anwendungen in Einfamilienhäusern mit Wärmepumpen, werden die Übertragbarkeit und Zuverlässigkeit der Methode unter unterschiedlichen Betriebsbedingungen prüfen. Darüber hinaus wird die direkte Integration zusätzlicher Kontextfaktoren — wie Wind, Solarstrahlung und Belegungsdynamik — in das Lernmodell sowohl die Prognosegenauigkeit als auch die Komfortregelung verbessern. Diese Entwicklungen werden den adaptiven Heizansatz als skalierbare und robuste Alternative zu statischen Regelstrategien stärken und den Übergang zu einem intelligenten, energieeffizienten Gebäudebetrieb vorantreiben.

Resumé

Cette recherche se concentre sur l'optimisation des systèmes de chauffage dans les bâtiments afin de réduire la consommation d'énergie et les émissions de CO_2. En Suisse, les bâtiments représentent une part importante de la consommation énergétique nationale, le chauffage en constituant la part dominante. Les systèmes de chauffage traditionnels reposent sur des courbes de chauffe statiques, configurées lors de l'installation, qui deviennent souvent obsolètes au fil du temps et entraînent des inefficacités systématiques. Ce projet introduit une approche adaptative qui ajuste continuellement la courbe de chauffe afin d'optimiser la consommation d'énergie tout en préservant le confort thermique.

L'objectif principal est d'ajuster la courbe de chauffe de manière autonome, sur la base d'un nombre minimal de données d'entrée — température intérieure, température extérieure et énergie de chauffage — sans recourir à un modèle explicite du bâtiment. Le système minimise une fonction de score combinant la consommation d'énergie normalisée et les écarts de confort. Des modèles de substitution par régression de processus gaussien (Gaussian Process) sont utilisés, par le biais de l'optimisation bayésienne contextuelle (Contextual Bayesian Optimization), pour identifier les paramètres optimaux de la courbe de chauffe au cours du temps. La température extérieure sert de variable contextuelle, ce qui permet à l'algorithme d'apprendre une stratégie dépendante du contexte plutôt qu'un point de consigne fixe unique.

Avant le déploiement sur le terrain, des simulations du système de chauffage ont été menées sur la base de modèles EnergyPlus détaillés pour Bülsweg — un immeuble résidentiel multifamilial — et NEST Sprint — une unité de recherche regroupant plusieurs bureaux. Ces simulations utilisent des données réelles du bâtiment, y compris des mesures issues de compteurs de chaleur et de capteurs, afin d'évaluer la consommation d'énergie de chauffage et le niveau de confort dans des conditions contrôlées, et servent à valider l'approche ainsi qu'à calibrer les hyperparamètres.

Des essais sur le terrain ont été effectués dans deux bâtiments aux typologies distinctes : Bülsweg (immeuble résidentiel multifamilial, en collaboration avec Lippuner AG, Buchs) et NEST Sprint (unité de bureaux, Empa Dübendorf). À Bülsweg, l'approche adaptative a permis de réaliser des économies d'énergie d'environ 4–6 % par rapport à la référence statique, tout en maintenant un niveau de confort pratiquement inchangé (Comfort Score : 0.24 → 0.23). Parmi les quatre méthodes d'évaluation appliquées — régression linéaire, normalisation par degrés-jours de chauffage (HDD), prédiction par apprentissage automatique et analyse par jumeau numérique — l'estimation basée sur les HDD s'est avérée la plus fiable, offrant un compromis équilibré entre normalisation météorologique et robustesse du modèle. Bien que le jumeau numérique n'ait pas pu reproduire ces gains, en raison d'une extrapolation limitée hors de sa plage d'apprentissage (30–45 °C), les données réelles ont confirmé des améliorations d'efficacité mesurables. Le modèle d'optimisation bayésienne a démontré une convergence stable après 30 à 40 itérations ; les tests de robustesse ont toutefois révélé une sensibilité asymétrique : tolérance face à de fortes valeurs aberrantes positives, mais vulnérabilité face aux valeurs négatives — soulignant la nécessité d'améliorer la gestion du bruit lors des futures itérations.

Pour NEST Sprint, l'évaluation a nécessité une correction saisonnière dédiée : la période de référence statique (octobre–novembre 2025) et la phase adaptative (février–mars 2026) tombent dans des saisons différentes, l'irradiance solaire en février–mars étant environ deux fois supérieure à celle d'octobre–novembre. Deux périodes de contrôle historiques — février–mars 2024 et 2025, toutes deux vérifiées comme étant en exploitation purement statique — ont été utilisées pour isoler le décalage saisonnier de la ligne de base de 10–12 % de l'effet de l'algorithme. Après élimination de ce facteur de confusion, l'économie nette imputable à AHA converge vers 5–10 % à travers deux méthodes indépendantes, avec une valeur centrale d'environ 7.5 %. Le confort thermique a été maintenu tout au long de la période, et la variabilité du confort a été réduite de 86 % (violation maximale du confort : 2.0 → 0.38). L'algorithme a démontré un comportement adaptatif au contexte statistiquement significatif, confirmant que le modèle de substitution GP a appris une stratégie de chauffage physiquement correcte. Un bug dans le calcul de la récompense, actif durant le déploiement, a été identifié a posteriori ;

une analyse contrefactuelle a démontré qu'il n'a eu qu'un impact marginal sur le choix des actions dans la plage de températures de fonctionnement, en raison de la structure dépendante du contexte de l'optimisation bayésienne.

Les travaux futurs porteront sur la mise à l'échelle et le perfectionnement du cadre de chauffage adaptatif afin de permettre une optimisation autonome et axée sur les données dans plusieurs bâtiments. La prochaine phase de développement vise à automatiser le traitement des données et le ré-entraînement des modèles, à renforcer la robustesse face au bruit et aux valeurs aberrantes des capteurs, et à élargir la plage d'optimisation pour exploiter pleinement le potentiel d'économie d'énergie du système. Un déploiement sur une saison complète dans les deux bâtiments est prévu pour l'hiver 2026/27, avec la résolution du couplage matériel Yc à Bülsweg et la mise en place d'un système automatisé de surveillance de l'état des capteurs. Des essais supplémentaires sur le terrain, y compris des applications dans des maisons individuelles équipées de pompes à chaleur, permettront d'évaluer la transférabilité et la fiabilité de la méthode dans des conditions d'exploitation variées. Par ailleurs, l'intégration directe de facteurs contextuels supplémentaires — tels que le vent, le rayonnement solaire et la dynamique d'occupation — dans le modèle d'apprentissage améliorera à la fois la précision prédictive et la régulation du confort. Ces développements consolideront l'approche de chauffage adaptatif comme une alternative évolutive et robuste aux stratégies de régulation statique, et feront progresser la transition vers un fonctionnement intelligent et économe en énergie des bâtiments.

Take-Home Message

The AHA framework offers a simple, interpretable, and transferable optimization concept. Its lightweight design and minimal data requirements make it adaptable to various heating systems, enabling scalable deployment across Switzerland's building stock. By improving control logic rather than hardware, the approach directly contributes to national goals of reducing building-related energy use and emissions.

Adaptive heating curves enable measurable energy savings without compromising comfort, as confirmed by field results at two buildings: 4–6% at Bülsweg (multi-family residential) and 5–10% at NEST Sprint (office unit), both with maintained thermal comfort.

Automation and robust data infrastructure are key to scalability. Manual data handling, unreliable sensors, and delayed system feedback highlighted the need for a fully automated and validated data pipeline to ensure reliable optimization across multiple buildings and heating systems.

Contents

List of abbreviations

SFOE	Definition: Swiss Federal Office of Energy
TB	Definition: Testbed. As Testbeds we can make use of NEST Umar, NEST SPRINT and Bülsweg.
DT	Definition: Digital Twin. When we refer to Digital Twins, we are actually referring to EnergyPlus models, which we use as a framework to simulate energy efficiencies. We use 'EnergyPlus models' and 'Digital Twins' interchangeably.
BO	Definition: Bayesian Optimization. We use Bayesian Optimization for a smart selection process to efficiently identify optimal parameters for the heating system.
CI	Definition: Confidence Interval quantifies the uncertainty associated with an estimated effect or parameter. In this work, the CI represents the 95% range within which the true value of the estimated energy savings is expected to lie, based on the variability observed in the data. The CIs were obtained using a bootstrapping procedure, which repeatedly resamples the data to construct an empirical distribution of the Difference-in-Differences estimates.
DID	Definition: The Difference-in-Differences method estimates the causal effect of an intervention by comparing the change in outcomes over time between a treatment and a control group. In this work, DID quantifies the impact of the adaptive heating tuner relative to the static baseline by evaluating the differential change in heat energy consumption predicted by two digital twins (DT1: adaptive, DT2: static). The approach is implemented in a bootstrapped framework to incorporate sampling variability and model uncertainty from the testbed.
GPR	Definition: Gaussian Process Regression. We use Gaussian Process Regression to approximate the underlying score function of heating energy consumption and comfort score, enabling a seamless implementation of Bayesian Optimization
Model	Definition: Our surrogate model, a Gaussian Process Regression model used within the Bayesian Optimization framework, learns the combined score (comfort + energy) as a function of inlet temperature (action) and context (24 hour mean outdoor temperature)
HDD	Definition: Heating Degree Days method, see Appendix for further explanations.
Inlet	Definition: Inlet water temperature of the heating circuit.
Action	Definition: In general, actions are the controllable parameter of a setting. In our specific implementation of a 1-dimensional linear heating curve action refers to the adjustment of the inlet temperature.
Context	Definition: In general, contexts are the non-controllable parameters of a setting e.g. weather conditions. In our specific implementation context refers to the mean outdoor temperature.
Energy Score	Definition: A normalized metric that quantifies the energy efficiency of the heating system over a 24-hour period, corrected for external weather influences such as sunshine duration and wind speed. Calculation: The raw heating energy consumption is denoised by subtracting the estimated influence of sunshine and wind, based on a linear regression model trained on historical data. This isolates the influence of outdoor temperature. The result is then normalized by dividing it by the expected energy consumption predicted by the denoising model for the given weather context. Interpretation: • A score of **1** means the adaptive heating curve performs around the estimated performance of the static configuration

- A score **< 1** indicates improved efficiency with respect to static configuration
- A score **> 1** indicates reduced efficiency with respect to static configuration

Formula:

$$f_{energy} = \frac{(Q_{heat} - \beta \cdot X_{weather})}{\alpha + \gamma \cdot X_{context}}$$

Q_{heat} = measured heating energy consumption

$X_{weather}$ = weather variables: sunshine duration and wind speed

$X_{context}$ = contextual variables, mean outdoor temperature

Comfort score

Definition: A penalty score that reflects how far the actual indoor temperature falls below a defined comfort threshold (20°C). It penalizes only underheating, as exceeding the comfort temperature is not considered a penalty because a heating system cannot cool when temperatures exceed the target.

Calculation: The score is computed using a smooth penalty function (log-sum-exp) applied only when the actual temperature is below the comfort threshold. This ensures that small deviations are tolerated, while large deviations are penalized increasingly.

Interpretation:

- A low score indicates good thermal comfort: temperature close to or above setpoint.
- A high score indicates poor comfort: temperature significantly below the threshold.

Formula:

$$f_{comfort} = \log(1 + \exp(T_{setpoint} - T_{actual}))$$

$T_{setpoint}$ = defined comfort threshold

T_{actual} = measured indoor temperature

1 Introduction

1.1 Context and motivation

Buildings account for approximately 42% of Switzerland's final energy consumption and 26% of total CO2 emissions, with heating demand being the primary driver at 68%. Heating systems are typically optimized using static heating curves set manually by technicians during installation. These initial settings are often suboptimal, degrade over time, and are rarely adjusted for changing conditions, such as tenant turnover. This leads to higher energy consumption, comfort issues, and increased scores.

Our adaptive heating curve tuning offers a continuous, self-adjusting solution. It relies on minimal data inputs – room temperature, heating and cooling energy use, outdoor temperature, and solar radiation – and optimizes using only 30 days of data. Daily updates require the building's total energy consumption and comfort deviations. This approach is lightweight, adaptable, and can run in the cloud or locally on the heating system's computer, ensuring efficient, responsive heating optimization.

1.2 Project objectives

The goal is to maintain thermal comfort while reducing energy consumption. The method autonomously adjusts the heating curve without altering the system structure. In this optimization problem, we aim to find the best heating curve parameters s to minimize our score function:

$$\min_{s \in S} \; y(\boldsymbol{s}, z) = a \cdot f_{energy}(\boldsymbol{s}, z) + b \cdot f_{Comfort}(\boldsymbol{s}, z) + \epsilon \tag{1.0}$$

- **Action s**: the heating curve parameters that control the heating system, in our setting the inlet temperature.
- **Context z**: the daily mean outdoor temperature, which influences heating energy use

We define a score as the value returned by the score function, which combines heating energy score and comfort score:

- **Heating energy score:** the function models the energy required for heating based on the parameters s and outside temperature z.
- **Comfort score:** the function measures how much the indoor temperature deviates from the desired setpoint.

The objective is to optimize s to minimize energy consumption while maintaining comfort, accounting for uncertainties through a noise term ϵ.

By incorporating Gaussian Process regression with an exponential kernel, the approach effectively models the relationship between heating energy, comfort, and ambient temperature. The method converges to the optimal heating curve settings, improving energy efficiency and comfort over time.

1.2.1.Comfort Score

To estimate comfort score, we calculate penalties for indoor temperatures that fall below the specified set point temperature. The calculation is performed separately for daytime (5:00–22:00) and nighttime (0:00–5:00 and 22:00–24:00) periods, recognizing their distinct comfort requirements. Temperatures above the setpoint are not penalized, and the total comfort score is normalized by the number of measurements taken during the respective period. For a clear illustration, refer to Figure 3 (comfort score function). The function we use to estimate the degree of daily comfort scores is as follows:

$$f_{comfort} = \log\left(1 + exp\left(T_{setpoint_{\frac{night}{day}}} - T_{actual}\right)\right) - \log(2) \tag{1.1}$$

This formula was chosen for its smoothness, as it avoids sharp transitions that could complicate model estimation of the underlying score function. Its design effectively penalizes deviations below the comfort

target while maintaining interpretability regarding the severity of comfort score. Small deviations from the set point result in minor penalties, whereas larger deviations are penalized exponentially, highlighting significant comfort score more prominently.

1.2.1.1. Comfort Objective: Comparison with ASHRAE

Fanger's model is the scientific standard for estimating indoor thermal comfort, using the metrics PMV (Predicted Mean Vote) and PPD (Predicted Percentage of Dissatisfied) to quantify how comfortable people feel in each environment. We did come up with our own comfort function due to simple reasons:

- PMV and PPD are nonlinear functions and can lead to problems while optimizing our objective function
- Fanger evaluates both temperature exceedances and deficits, which is problematic in our context since a heating system cannot cool when temperatures exceed the target, and such cases should therefore not be penalized.

In contrast, our comfort violation formula provides a continuous and more efficient approach, focusing on the key variable temperature, making it better suited for optimization problems.

A direct comparison of PMV and PPD for some selected days during our heating period display that is with respect to correlation highly similar. The Base assumption for comparison is a standard case in winter: activity = "Typing", garments = ["Sweatpants", "Long sleeve shirt (thick)", "Thick trousers", "Calf length socks", "Slippers"]

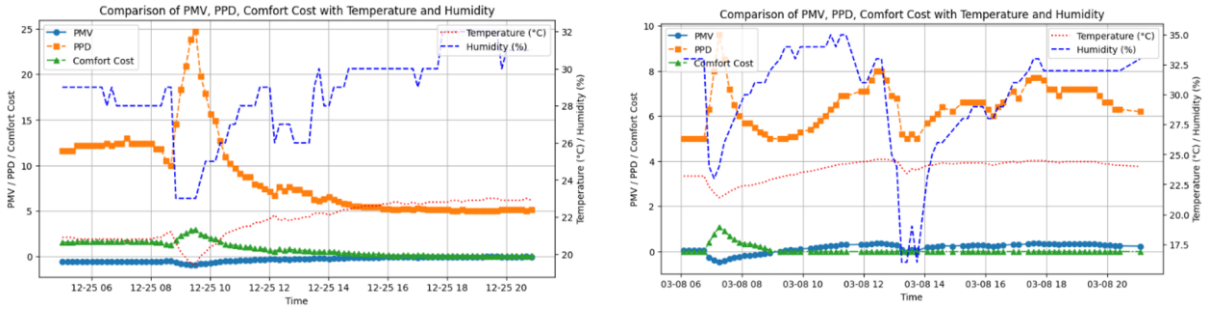

Figure 2: Displayed is a direct comparison of PMV, PPD, and our comfort score metric over the course of the selected days: 25.12.24 and 08.03.2025. Aside from the obvious differences in scale, the overall behaviour is notably similar.

For these two days, we calculated the Pearson correlation coefficients and obtained the following results:

- PMV and Comfort Score: Very strong negative correlation of -0.96
- PPD and Comfort Score: Very strong positive correlation of +0.98

These high correlations indicate that both PMV and PPD are strongly aligned with our comfort score metric. This suggests that the proposed comfort score approximation captures the essential information of traditional comfort indicators and is therefore an efficient candidate for estimating occupant comfort.

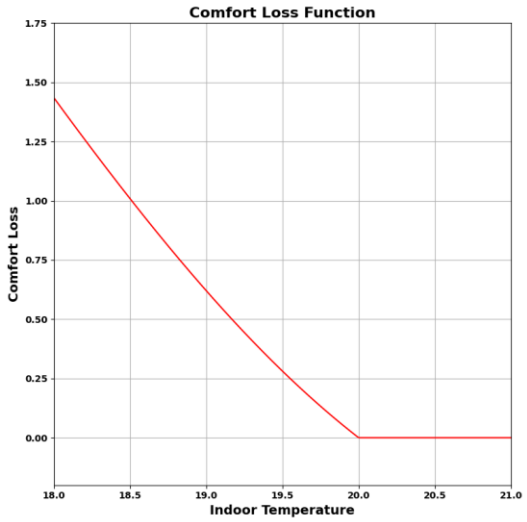

Figure 3: Shows comfort score estimated by our comfort loss function. We only consider comfort scores which are temperatures below a defined threshold, because the heating system can only warm but not cool.

2 Approach and Method

In our approach a simple two-point parametrization of the heating curve is proposed together with the adaptive heating curve strategy GPR and BO, see for more details appendix 5.1. The adaptive solutions are discussed regarding the tracking of the defined score based on the desired indoor temperature $T_{setpoint}$ and heating energy in simulations and field tests.

2.1 Parameterization of the heating curve

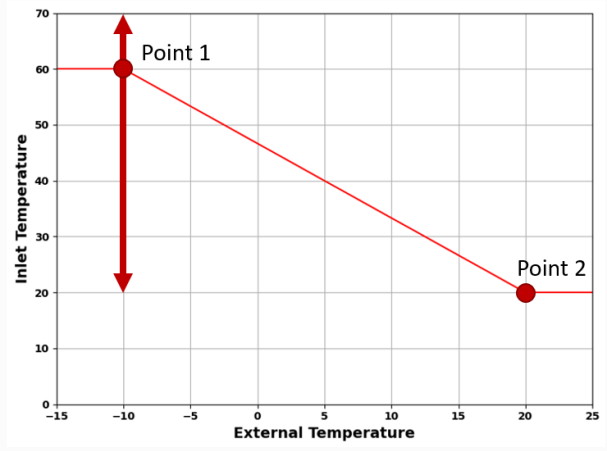

Figure 4: Shown is a 2 point linear heating curve defined by Point 1 and Point 2. We focus on a 1-dimensional linear heating curve, specifically on reference point 1 on $T1_{inlet}$, which represents the inlet temperature. Highlighted for reference is the potential safe range of parameter values considered during our Bayesian optimization process, ranging from 20 up to 70 degrees Celsius. This ensures that the optimization remains within safe operational limits while seeking the optimal solution. Right:

The heating curve defines the relationship between the outdoor temperature and the inlet temperature of the heating system, ensuring efficient, weather-controlled operation. For heat pump-based floor

heating, optimizing the heating curve helps maintain thermal comfort, reduce operational scores, and maximize the heat pump's efficiency. Typically, the heating curve is linear, as heating demand is approximately proportional to outdoor temperature [5].

We simplified our optimization to a one-dimensional linear heating curve, see figure 4, parameterized using two reference points based on the mean external temperature:

- Reference Point 1 = $(T1_{external}, T1_{inlet})$
- Reference Point 2 = $(T2_{external}, T2_{inlet})$

This approach leads to a strong simplification, ensuring easier interpretation, reduced data requirements, and straightforward adaptation for optimal system performance.

2.2 Adaptive optimization of the heating curve

Our adaptive heating strategy automatically adjusts the heating curve daily, without requiring detailed prior modelling of the building or heating system. To do this, we define a generic score function (1.0), which is initially estimated for each new building and then gradually optimized over time. This function calculates a score based on thermal comfort score and energy score, which depend on the selected action – the inlet temperature $T1_{inlet}$ – and the prevailing context – the mean outdoor temperature. Intuitively, certain combinations of inlet and outdoor temperatures will result in a higher (i.e., more scorely) score – for example, high inlet temperatures during warm outdoor conditions, or very low inlet temperatures when it's cold. In contrast, high inlet temperatures during cold weather are more reasonable and should lead to lower scores.

The exact behaviour and magnitude of this score are building-specific and must be learned through continuous sampling and evaluation of the relationship between inlet temperature, context, and resulting energy and comfort scores which defines the score of a building. To guide the sampling process, we define a “safe space” – a range of inlet temperatures that are considered safe for the given building and heating system. For example, for underfloor heating, an inlet range of 25 - 45°C is typically safe. We begin by sampling 10 predefined points from this safe space. Example for Bülsweg: [30–45°C].

After this initial exploration, the actual optimization begins. The data collected from the initial samples is used to fit a first version of the scoring function. This fitted model can then be used to determine the optimal inlet temperature for a given context – specifically, the estimated outdoor temperature for the next day – to minimize the score.

This process of sampling, updating the model, and selecting the best possible inlet temperature is repeated iteratively each day.

In detail, the procedure for adapting and optimizing the algorithm works as follows: We use Bayesian Optimization (BO) to guide this process. BO systematically predicts based on the latest measurement the next best $T1_{inlet}$ most likely to reduce the score function. Each measurement contributes a new data point to evaluate the score function, which is modelled using GPR. This approach enables efficient exploration of the available parameter space of the heating curve.

The adaption of the reference point 1 based on the mean outdoor temperature is performed according to the problem iteration cycle [2, 3]:

2.2.1.Problem Iteration Cycle

Pseudocode of our contextual Bayesian Optimization approach to optimize the heating curve, see also figure 5 for a summarized description:

1. Estimation of Context: mean outdoor temperature of previous day:

We take contextual information into account only above a certain threshold of samples, hard coded 10, and focus initially solely on a safe space for $T1_{inlet}$ itself. This space of safe points contains 10 entries ranging from [30–45°C] for Bülsweg.

2. **Estimation of next $T1_{inlet}$ based on context from 1 and existing surrogate model with Bayesian Optimization algorithm, upper confidence bound.**

We start to use upper confidence bound only after a list of safe points have been sampled first. BO using the upper confidence bound strategy uses the updated surrogate model to make it's prediction of the next optimal inlet temperature.

3. **Update heating curve parameters with next $T1_{inlet}$ and start simulation**

Energyplus engine creates an energyplusoutput file which is processed by Simulationdriver class to be used for the ML process

4. **The heating energy from the Energyplusmodel is processed by our Denoiser to denoise and normalize the heating energy. This result is then combined with the estimated comfort violation to estimate our score, which represents our new sample.**
5. **The new sample is used to update our GPR**
6. **The iteration cycles starts over, by estimating the new context via mean temperature of the previous day**

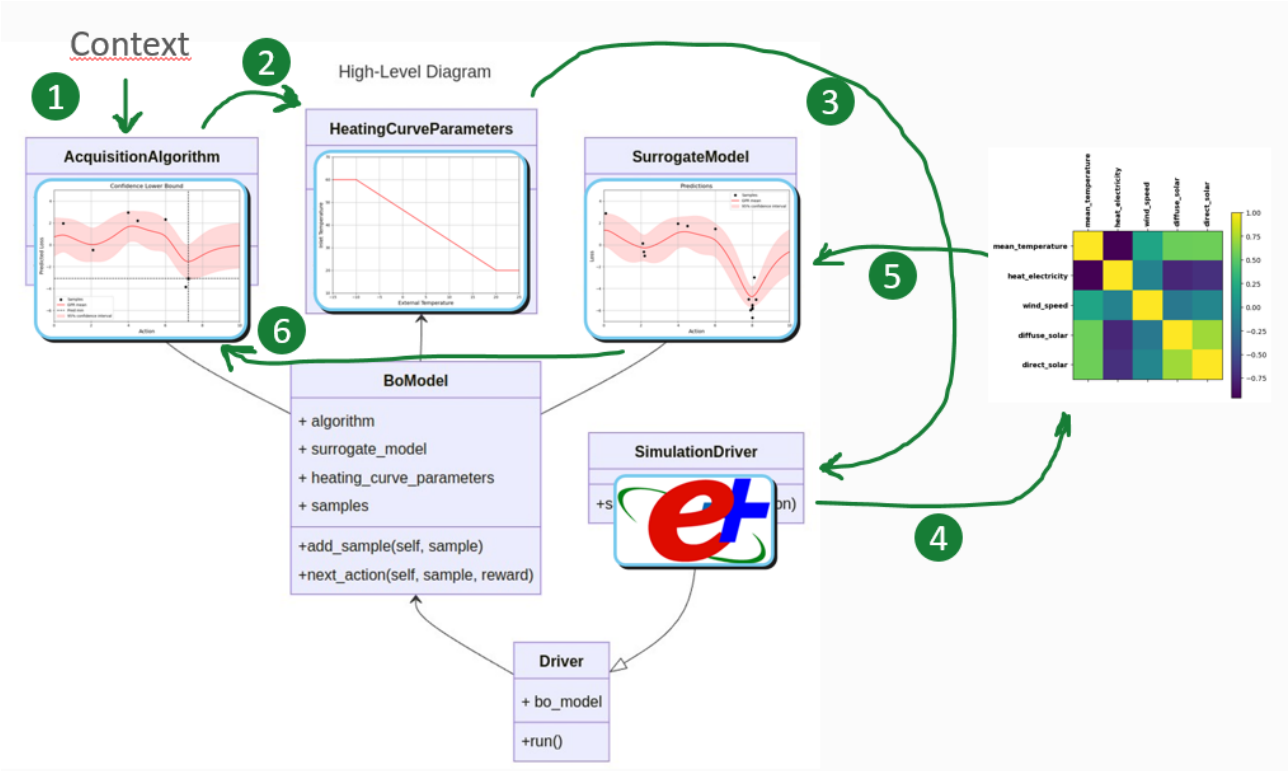

Figure 5: The process begins with the estimation of the context, using the mean temperature of the previous day. Next, the algorithm estimates the next action using the context and an existing surrogate model, applying the Upper Confidence Bound method to obtain the new heating parameters. With the heating curve parameters updated the EnergyPlus simulation is run to generate heating energy and comfort violation data. The heating energy data is denoised, normalized and combined with comfort score to estimate the score. This new sample is used to update the Gaussian Process regressor – which mimics our understanding of the score function, and the cycle repeats by estimating the new context.

3 Results and Discussion

To demonstrate the performance of the proposed adaptive strategy and validate our claim that a dynamic, adaptive heating tuner can achieve a 5% reduction in energy consumption while maintaining comfort, the following objectives must be shown:

- Energy Efficiency: Achieve a 5% reduction in energy consumption while maintaining comfort levels throughout the entire heating period.
- Convergence Behaviour: After 30 iterations of optimizing the heating curve per context, there should be no significant changes in regret, difference between an optimal.
- Transferability: The algorithm must be applicable across different building types and heating systems, delivering comparable results.
- Robustness: The system should effectively respond to extreme climate and user scenarios and handle missing data without performance degradation.

To achieve these goals, we will conduct simulations followed by real-world field tests. The simulations are based on detailed EnergyPlus models to replicate various building and system configurations. For further details on the simulation models, see chapter 7.2. Some of the objectives listed above will be analysed only within the simulations: convergence and robustness. To analyse convergence properly, we need an idea of an optimal curve, which can be found through grid search using contextual data. This will allow us to specifically test whether our adaptive strategy converges toward this curve. We will test robustness by applying sudden temperature changes, which are rare in real-world scenarios and difficult to predict.

3.1 Simulation

3.1.1. Energy Score

To assess the normalized energy consumption – the energy score, see problem iteration cycle 2.2.1. point 4 – we fixed the inlet temperature $T1_{inlet}$ to a specific value e.g., 30°C for one heating period and simulated the entire period for NEST, see figure 6. This process was repeated for inlet temperatures ranging from 30°C to 60°C. Each data point for a given inlet temperature corresponds to the total daily heating energy consumption. Additionally, each day is color-coded based on the average outdoor temperature: **blue** indicates colder conditions (around 0 °C), while **red** indicates warmer conditions (around 8 °C).

It is evident that the Energy score is lowest for the lowest inlet temperature and increases as the inlet temperature rises. Another observable trend is that the increase in scores levels off above 50°C, primarily because higher inlet temperatures reduce the time required to heat the space to the desired set point temperature.

What the Energy Score does not reveal, however, is that at inlet temperatures of 30°C or lower, comfort is often significantly compromised, see next chapter.

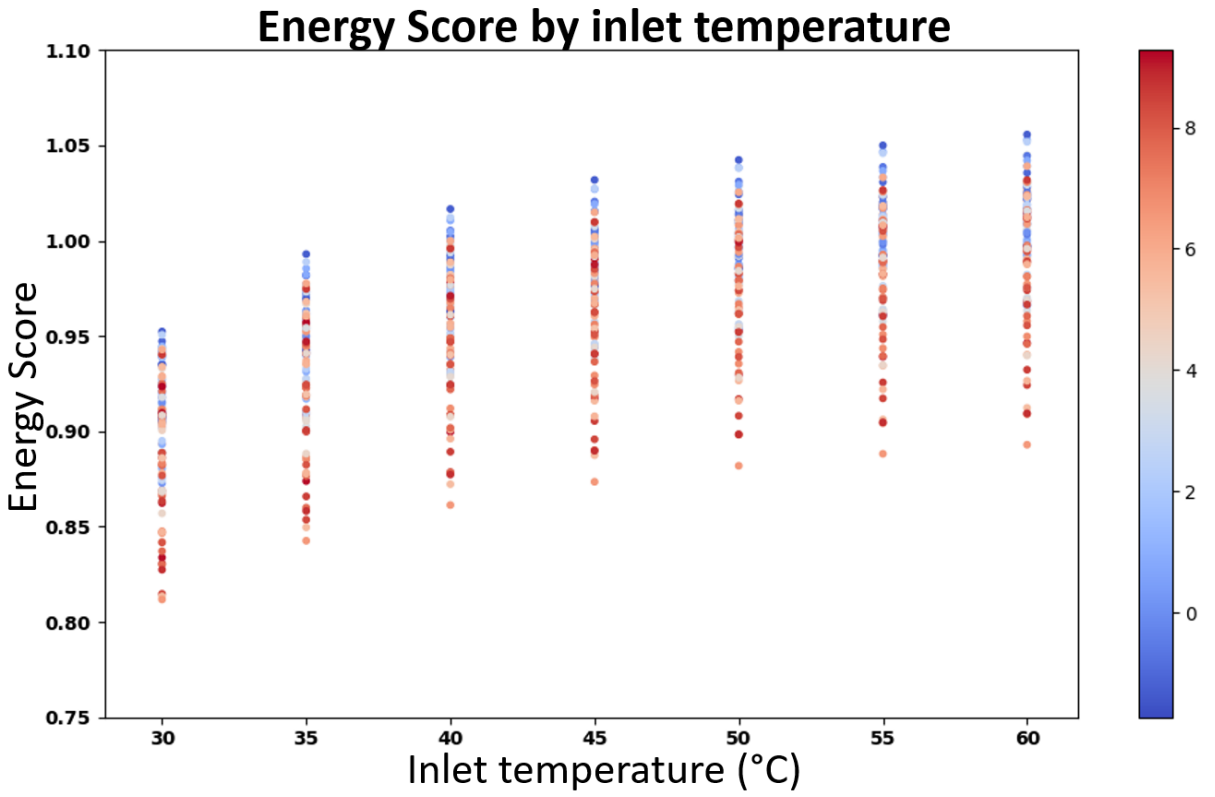

Figure 6: We compute the total consumed heating energy remove noise and normalize the heating energy consumption to a score independent of external conditions, as described in section 2.2.1, point 4. Normalization is based on the default heating curve, where the consumed energy is divided by the energy consumption of the default heating curve. Because the default heating curve is defined at $(T1_{external}, T1_{inlet}) = (-10, 36)$ and $(T2_{external}, T2_{inlet}) = (20, 20)$, heating curves with a lower T_{inlet} will have a reduced heating energy score.

3.1.2.Comfort Score

To assess the comfort score, see 1.2.1, we fixed the inlet temperature to a specific value (e.g., 30°C) for one heating period and simulated the entire period for NEST, see figure 7. This process was repeated for inlet temperatures ranging from 30°C to 60°C. Each data point for a given inlet temperature corresponds to the total daily consumption. Additionally, individual days are color-coded based on context – specifically, the average outdoor temperature: blue represents an average temperature around 0°C, and red represents an average temperature around 5°C.

We can clearly observe that at inlet temperatures above 40°C, there are virtually no comfort scores. Below 40°C, however, significant comfort scores occur, particularly on colder days, while on warmer days, comfort scores are largely absent. These comfort scores arise primarily for two reasons: first, the target temperature at cold outdoor temperature cannot longer be reached with lower inlet temperatures, and second, it takes longer to heat the building to the desired temperature.

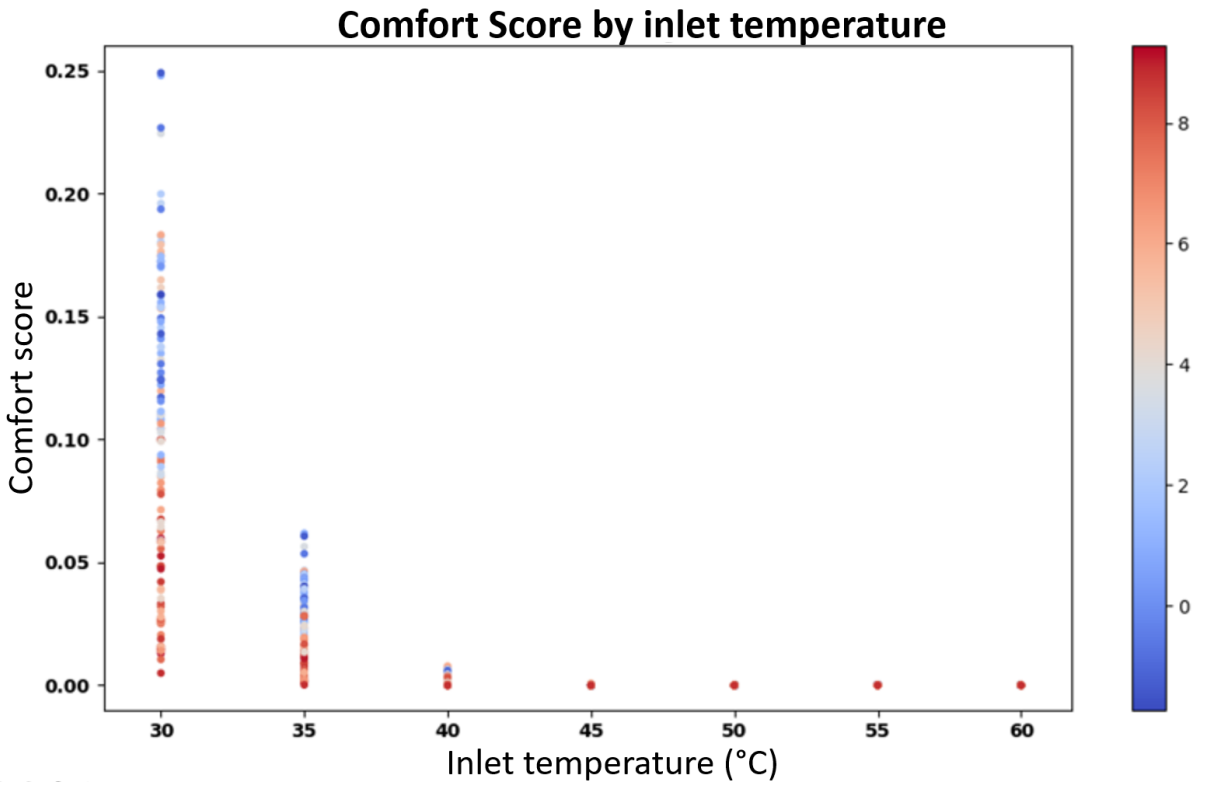

Figure 7: Comfort score for varying inlet temperatures (30°C–60°C). Daily energy consumption is plotted with color coding based on the average outdoor temperature: blue for 0°C and red for 5°C. Comfort score are minimal above 40°C but increase below, particularly on colder days, due to slower heating and difficulty reaching the target temperature.

3.1.3.Combined Score and Score Function

We assumed a static heating curve and varied the inlet temperature between 25°C and 60°C, calculating the resulting scores as the sum of normalized heating energy and daily comfort score, combined score. This analysis was performed for all days within the selected heating period, with each individual point representing a single day, see figure 8 on the left.

Using GPR, we estimated the score function for each context, resulting in a series of curves. These curves are color-coded based on their context: dark blue represents an average outdoor temperature of -2°C, while dark red corresponds to 10°C. The score functions enable us to identify the optimal inlet temperature settings for the heating curves. For instance, the minimum of the dark blue curve is around 40°C, see figure 8 on the right.

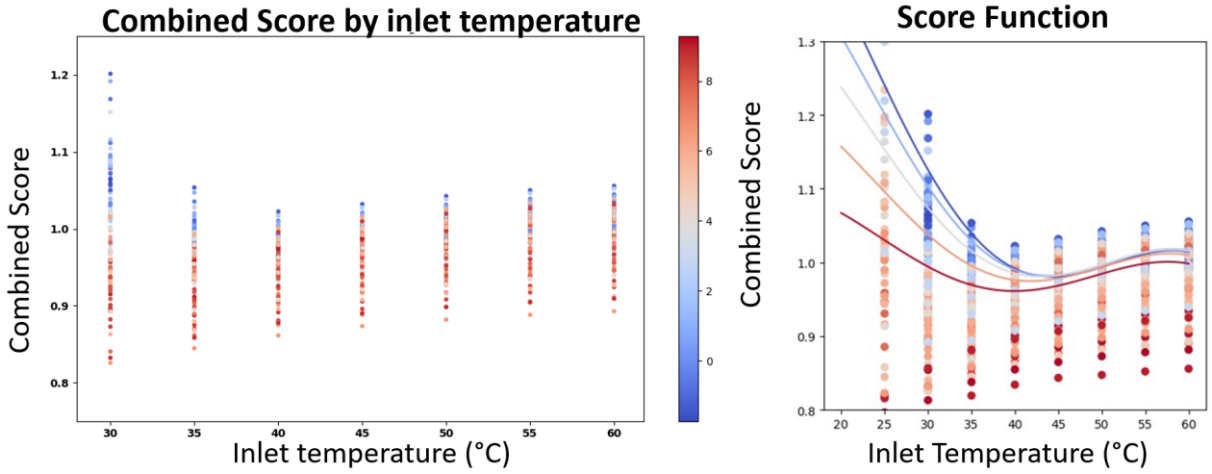

Figure 8: On the left, we show the combined score based on our score function from Chapter 1.2, which is a weighted linear combination. The weights, set as trade-off parameters, significantly impact the potential energy savings, see Appendix chapter 5.4. Here, a weight of 1 is applied, prioritizing low comfort score, especially at lower temperatures, which limits the energy-saving potential. On the right, the trained model made predictions based on context and inlet temperatures. The obtained curves clearly illustrate a distinct minimum, indicating the optimal inlet temperature.

3.1.4.Results Bülsweg and NEST

The adaptive results were calculated using the proposed strategy – the Problem Iteration Cycle 2.2.1. Key findings are illustrated in figures 9, 10 and summarized in the following subsection:

The results of the adaptive heating strategy for NEST and Bülsweg are similar regarding the development of the score function. Thus, selected results for NEST are presented, comparing the adaptive heating curve with two static heating curves with reference points: ($T1_{external}$, $T1_{inlet}$) = $(-10, 60)$, and $(-10, 40)$. We conducted three simulations, one for each heating curve strategy: adaptive, high and medium static heating curves. The daily scores were calculated and plotted for each heating curve to enable a direct comparison. The adaptive heating curve shows that score reductions can be measured as early as the third day, and the adaptive strategy outperforms both static curves. In this example, the adaptive heating strategy achieves energy savings of approximately 1.5% for the chosen heating period of January – March 2021 compared to the lower curve and ~5.9% compared to the higher curve, quickly converging to the optimal solution for each mean outdoor temperature.

Figure 10 presents the results for the best possible choice of $\mathrm{T1_{inlet}}$ based on the context and outdoor temperature. The context is color-coded: dark blue corresponds to an average outdoor temperature of -2°C, and dark red to 10°C, compare chapter 3.1.3. We marked the minimum point for each curve. For example, the estimated score function for NEST - Umar and Bülsweg indicates that the minimum for average daily temperatures of -2°C occurs at an inlet temperature of 40°C.

A direct comparison with the initial heating curve values for Bülsweg: ($T1_{external}$,$T1_{inlet}$) = $(-10, 36)$ and ($T2_{external}$,$T2_{inlet}$) = $(20, 20)$ reveals that the heating curve is only optimal for an average outdoor temperature of 0°C. In contexts deviating from this temperature, score reductions can be achieved by adapting $T1_{inlet}$, with particularly significant savings in warmer contexts. For NEST - Umar, a similar pattern is observed.

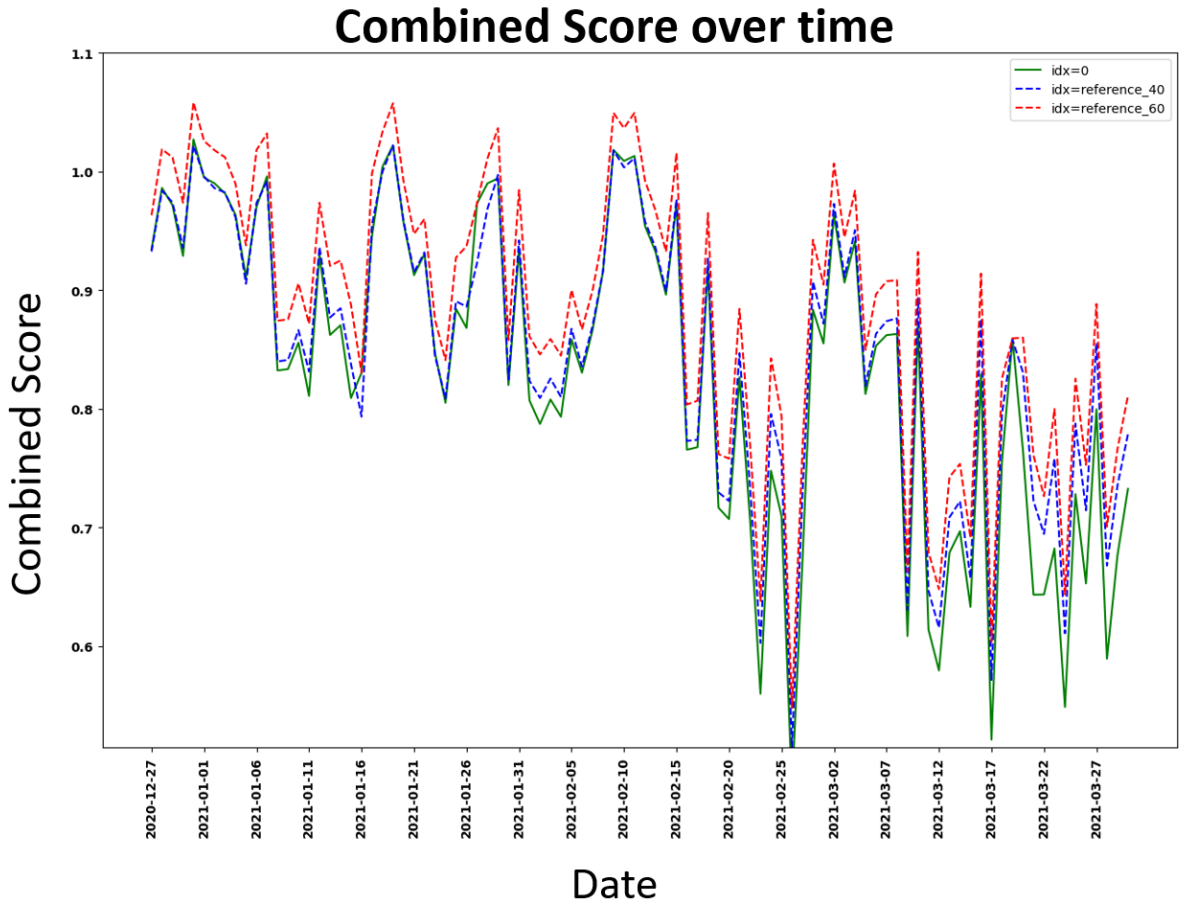

Figure 9: The results for a 3 month simulation of NEST Umar with our adapted heating curve, idx=0, and two reference heating curves fixed at an inlet temperature at 40 and 60 degrees.

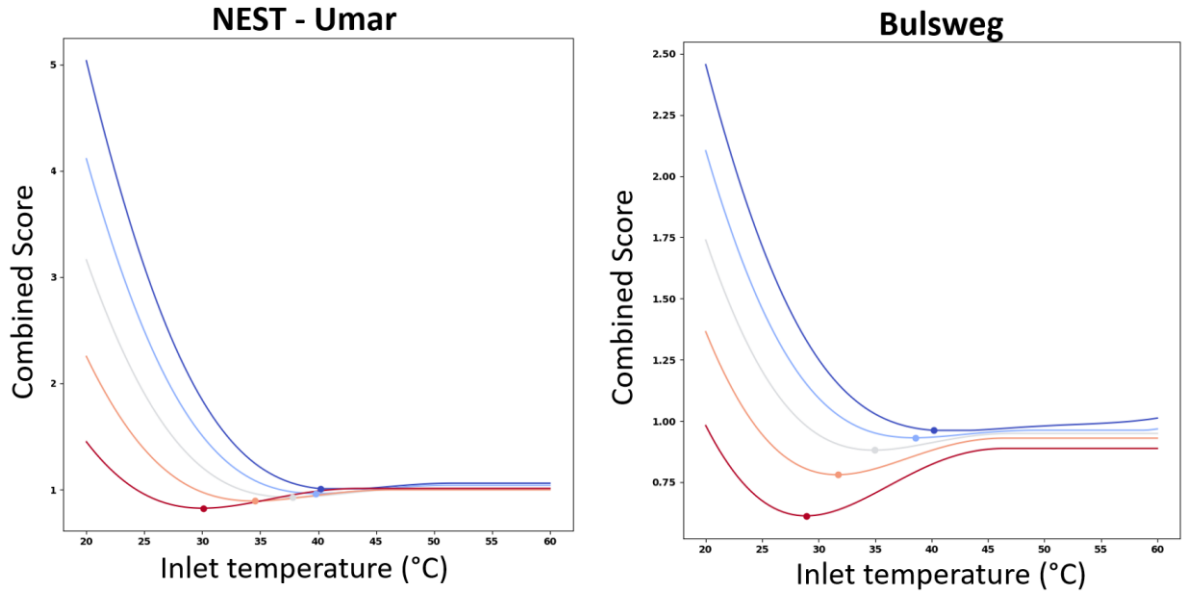

Figure 10: Shows score functions obtained by GPR. A direct comparison of the two DT for NEST resp. Bülsweg. Highlighted are 5 different graphs corresponding to the context: -2, 0, 2.5, 6.25, 10 degrees Celsius from dark blue to dark red, based on the surrogate estimation of the score per action.

3.2 Field Test / Benchmark

To assess AHA under real building conditions, field tests were conducted in two buildings with distinct typologies and heating systems. The first test site is Bülsweg, a multi-family residential building in Buchs, Switzerland, provided by our industrial partner Lippuner AG (see Appendix 8). The building comprises nine apartments (three 2.5-room and six 3.5-room units) connected to the district heating network of the Buchs waste incineration plant, with underfloor heating throughout. Heat consumption per apartment is monitored via digital heat meters. Indoor environmental quality is measured via multi-function sensors (VOC, temperature, relative humidity). Meteorological data is sourced from the Vaduz weather station. The second test site is NEST Sprint, a 167 m² office unit comprising 12 individual offices at the NEST research platform at Empa Dübendorf, opened in August 2021. The unit uses a ceiling heating/cooling system connected to the NEST Medium Temperature Energy grid via a heat exchanger. Room temperatures are individually controllable via one valve per room. The supply temperature is governed by a weather-compensated heating curve — the parameter optimized by AHA. The field test results are presented in two parts: Section 3.2.2 covers Bülsweg; Section 3.2.3 presents NEST Sprint. Section 3.4 discusses transferability across the two building typologies.

Figure 11 illustrates the dynamic adjustment of the heating curve of Bülsweg in response to outdoor temperature variations. The subsequent field tests aim to quantify whether this adaptive strategy reduces energy consumption compared to the conventional static heating curve.

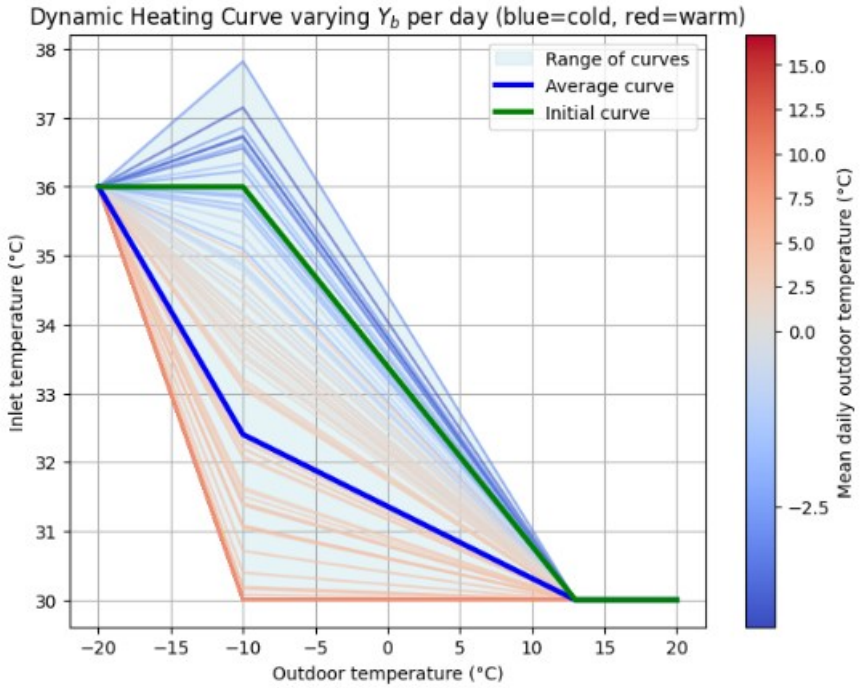

Figure 11: Dynamic heating curves generated by Bayesian Optimization. The green line shows the static baseline, the blue line the average BO-adjusted curve, and individual daily curves are colored by mean daily outdoor temperature (blue = cold, red = warm), as indicated by the colorbar. The light blue shaded area shows the full range of curves. On average, the BO curve lowers the inlet temperature to ~32.4 °C at –10 °C outdoor temperature (~31 °C at 0 °C). During November – December 2024 and some days in January 2025, ~63% of curves remain below 32 °C, while occasional upward adjustments reach up to 37.8 °C to maintain comfort. The figure illus-

trates that for warmer outside temperatures, BO tends to lower the inlet temperature, while for colder temperatures, it occasionally increases it. Overall, this adaptive adjustment suggests potential energy savings by reducing heating demand while still ensuring comfort.

3.2.1. Benchmark Analysis

Our study employs a "single-group pre-post design": measurements are taken before and after the intervention within the same building to capture changes attributable to the adaptive heating algorithm. We assume that user behaviour is not significantly affected by the intervention and that energy consumption is primarily driven by context variables: outdoor temperature and solar irradiance.

For both test sites, the static heating curve serves as the pre-intervention baseline; the AHA-adaptive phase constitutes the intervention. The evaluation approaches differ between the two sites due to a key structural difference: at Bülsweg, the static reference period and the adaptive phase overlap substantially in season, making a direct ML-based virtual control group feasible. At NEST Sprint, the two periods fall in different seasons (October–November vs. February–March), which introduces a structural solar confounder that renders ML-based extrapolation unreliable. For Sprint, two historical control periods (February–March 2024 and 2025, both without AHA, non-active experiments) were used instead to isolate the seasonal baseline shift from the algorithm effect.

The following table summarizes the evaluation methods applied at each site and their principal results.

Analysis	Bülsweg Savings	Sprint Savings
Regression (E ~ T)	5.7%	17.6%
HDD normalization	4.1% ± 1.1%	22.4% ± 0.7% (unadjusted)
ML virtual control group	1.6%	—[1]
Digital Twin DID	not significant [2]	—[1]
Control period comparison	—	8.8 − 11.5%
Multivariate weather correction	—	5 – 10%
Comfort	maintained (0.23 vs. 0.24)	maintained
Context–Action correlation	confirmed	r = −0.537, p = 0.003

Interpretation. For Bülsweg, the HDD-normalized estimate of 4.1% is the most defensible figure: it corrects for temperature variability without strong model assumptions, and is bracketed by the regression estimate (5.7%) and the ML virtual control group (1.6%). For Sprint, the unadjusted figures 17.6% - 22.4 % overstate the AHA contribution due to the seasonal confounder. Once the 10–12% seasonal baseline shift is removed via the control period comparison, the net AHA-attributable saving converges to 5–10%. In both buildings, thermal comfort was maintained throughout the adaptive phase.

Our algorithm optimizes an objective function combining normalized heating energy and thermal comfort into a single dimensionless reward (see Section 1.1 and 3.1). To better interpret the results, the individual effects on energy use and comfort are presented separately in the following sections.

The field test results are presented in two parts: Section 3.2.2 covers the Bülsweg multi-family residential building, including the full description of all evaluation methods. Section 3.2.3 presents the results

[1] Not applied: static reference (Oct–Nov) and adaptive phase (Feb–Mar) fall in different seasons; ML extrapolation would conflate the AHA effect with a seasonal baseline shift of 10–12%.
[2] 95% CI: −39.7 to +17.5 kWh/day

for NEST Sprint, referencing the methodological descriptions from Section 3.2.2 where applicable and introducing the control period comparison as a Sprint-specific extension.

Our algorithm optimizes the objective function 1.0, which combines normalized heating energy and comfort score into a single, dimensionless score. To better interpret the results, we present the individual effects on energy use and comfort.

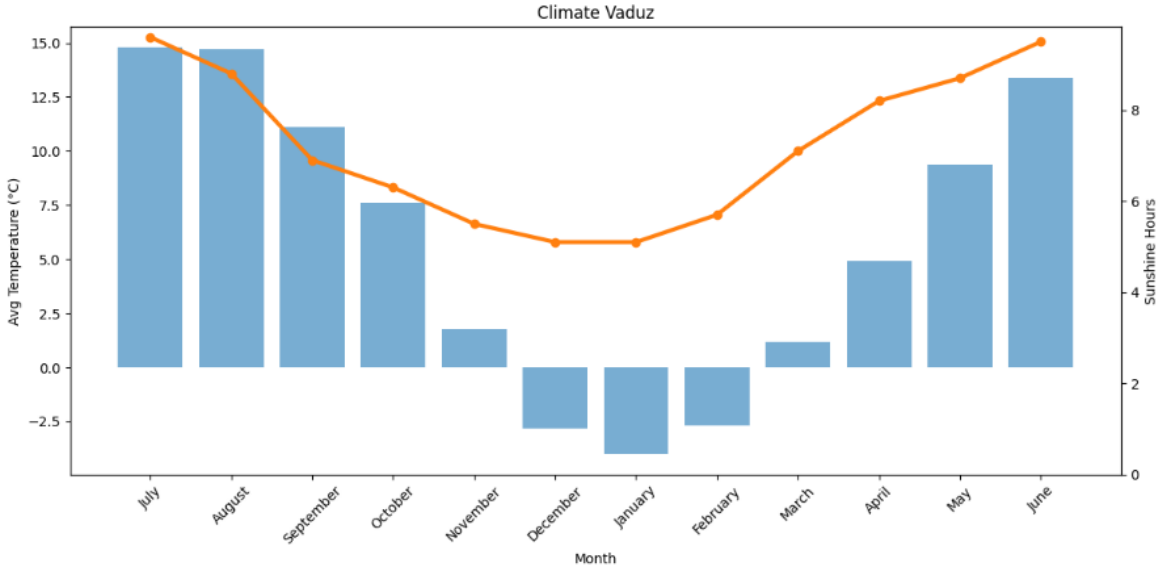

Figure 12: The image shows Vaduz's climate diagram for the climate reference period 1981−2010. It clearly illustrates the symmetry from January onward between the first and second halves of winter concerning major external influencing factors: hours of sunlight and temperature.

3.2.2. Bülsweg

For Bülsweg, the intervention was planned to maximize days with similar temperature and sunshine profiles. Validation of the digital twins occurred before the intervention, with accuracy improving over longer validation periods. A critical metric is convergence: the goal is to complete the training phase within 30 days (upper limit for the heating curve model), though a shorter period was anticipated (see Section 3.1). Days required for convergence were excluded from the comparable sample; a continuous intervention was therefore preferred. Based on the symmetrical climate chart for Vaduz (Figure 12), January 2025 was selected as the intervention period.

Measurements at Bülsweg cover the heating season from November 15, 2024, to March 31, 2025. The adaptive heating curve was activated on January 23, 2025, and remained active for two months. Outdoor temperatures during this period ranged from −2°C to 15°C. To allow the GPR model to approximate the score function over the action-context space, a minimum of 10 data points was required. The first 10 samples were therefore excluded from the analysis: a 5-day safe exploration of predefined inlet temperatures [30–45°C] followed by 5 days of additional sampling. Further details on data and recording frequencies are in Appendix 8.

3.2.2.1. Regression Analysis

We will begin by comparing the data before and after the intervention using linear regression. Two linear curves will represent the average heating scores against the average daily temperature, pre- and post-intervention.

Before intervention:

$$y_{static} = m_{static} * x + b_{static} \quad (3.0)$$

After intervention:

$$y_{adaptive} = m_{adaptive} * x + b_{adaptive} \quad (3.1)$$

Where: x = average daily outdoor temperature and y = objective score (heating energy plus comfort score) for the observed day.

The difference between these two lines tells us how much energy we save by using the adaptive algorithm at a specific outdoor temperature x

$$\Delta_{Effect} = \Delta_{slope} * x + \Delta_b \quad (3.2)$$

$$\Delta_{slope} = m_{static} - m_{adaptive}$$

$$\Delta_b = b_{static} - b_{adaptive}$$

This provides an efficient estimation of the expected effect of the algorithm over a heating period [4].

We begin by focusing on the heating energy component: From figure 13 we can see that the static heating configuration is less efficient overall but reacts more aggressively to warmer outdoor temperatures leading to larger score drops in warm conditions. The adaptive strategy is more stable: saving more in cold conditions, possibly less in warm ones. Comparing the two regression we get:

$$y_{static} = -8.1366 * x + 201.54$$

$$y_{adaptive} = -4.1578 * x + 167.14$$

$$\Delta_{slope} = -3.9789$$

$$\Delta_b = 34.40$$

From these results we can calculate the weighted energy savings by context distribution:

$$\Delta Energy_{weighted} = \sum_{x \in context} p(x) \cdot (\Delta_{slope} \cdot x + \Delta_b)$$

$p(x) = \frac{n_x}{N}$ = relative weight (in %)

n_x = number of observations per context x

N = total number of observations

and obtain an estimated energy saving of 146 kWh over the temperature range during the active time of the adaptive heating system and an average weighted energy saving per degree Celsius of 11 kWh. The total estimated heating energy consumption over the temperature range from −2°C to 15°C for the static configuration amounts to approximately 2527 kWh. Based on our regression-based estimate of 146 kWh in savings, this corresponds to an energy reduction potential of about 5.7% (= $\frac{146\ \text{kWh}}{2527\ kWh}$) over the evaluated range.

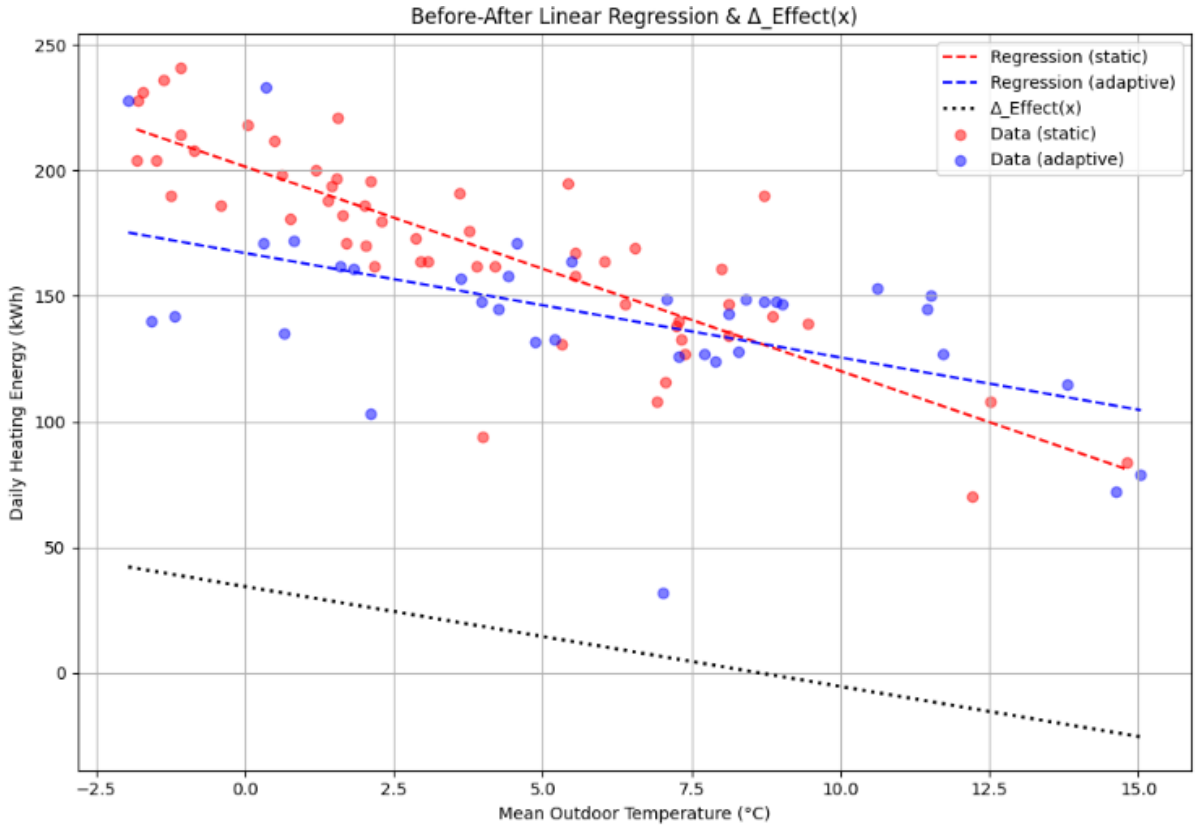

Figure 13: The scatterplot displays the analyzed data points for Bülsweg, including the fitted regression lines for the static (red), adaptive (blue), and net effect (black) configurations. To ensure a fair comparison and minimize extrapolation errors, we only included data points from the pre- and post-intervention periods that fall within the temperature range where the adaptive heating curve was active. Additionally, the first 10 samples from the adaptive configuration were excluded to account for the initial exploration phase of the optimization. The net effect regression suggests an estimated total energy saving of approximately 146 kWh across the explored temperature range, corresponding to about 11 kWh per degree Celsius.

3.2.2.2. Heating degree day method:

To estimate the net energy-saving effect of the new adaptive heating control implementation, we used a HDD normalization approach. This method corrects for differences in outdoor temperature by quantifying heating demand using the HDD metric:

$$HDD = \sum_{k=0}^{n} \max(0, \mathrm{T_{set}} - \mathrm{T_{outside}})$$

Based on this, we calculated the relative change in specific energy consumption per HDD before and after the intervention, see for more details Appendix 5.2:

$$\Delta_{Effect} = \left[1 - \frac{\frac{\mathrm{Heatingenergy_{post}}}{HDD_{post}}}{\frac{\mathrm{Heatingenergy_{pre}}}{HDD_{pre}}}\right] \cdot 100$$

$$\mathrm{Heatingenergy} = \sum_{k=0}^{n} Heatingenergy_k$$

This approach yielded an estimated energy-saving potential of approximately 4.1±1.1%.

3.2.2.3. Machine Learning:

We trained machine learning models on pre-intervention data to predict the heat energy consumption in kWh. Since detailed occupant information (e.g., presence, window usage, or shading behaviour) is not directly available, we infer these effects from the historical data. We assume that such behavioural patterns remain consistent over time and are unaffected by the intervention.

We used weather features as inputs to three models: a Generalized Additive Model, our BO and XGBoost. Each model was trained using data from September 2023 to March 2024 and evaluated on a hold-out test set. The XGBoost model, optimized with Optuna for hyperparameter tuning, achieved the best performance with an average RMSE of 15.35 and R^2 of 0.89 based on 18 optimization runs with R^2 values ranging from 0.884 to 0.903 and RMSE values between 14.46 and 15.83.

To quantify the algorithm's impact, we compared predicted consumption with actual post-intervention measurements. The difference estimates the net energy saving:

$$\Delta_{Effect} = Heatingenergy_k - \mathrm{Prediction_k}$$

To further assess the energy-saving potential, we performed a regression analysis on both the model predictions and the actual measurements as functions of outdoor temperature. Linear regression lines were fitted to the predicted baseline and the measured consumption data. We then integrated the area under each regression line over the temperature range from −2°C to 15°C, which represents the active operating range of the adaptive heating system. The difference between these areas provides an estimate of the total energy savings. The XGBoost model estimates an average consumption of 2418 kWh. Based on our regression analysis, the intervention led to an estimated average energy saving of 40 kWh, corresponding to an average reduction potential of approximately 1.6% (= $\frac{40\ \mathrm{kWh}}{2418\ kWh}$) over the evaluated range.

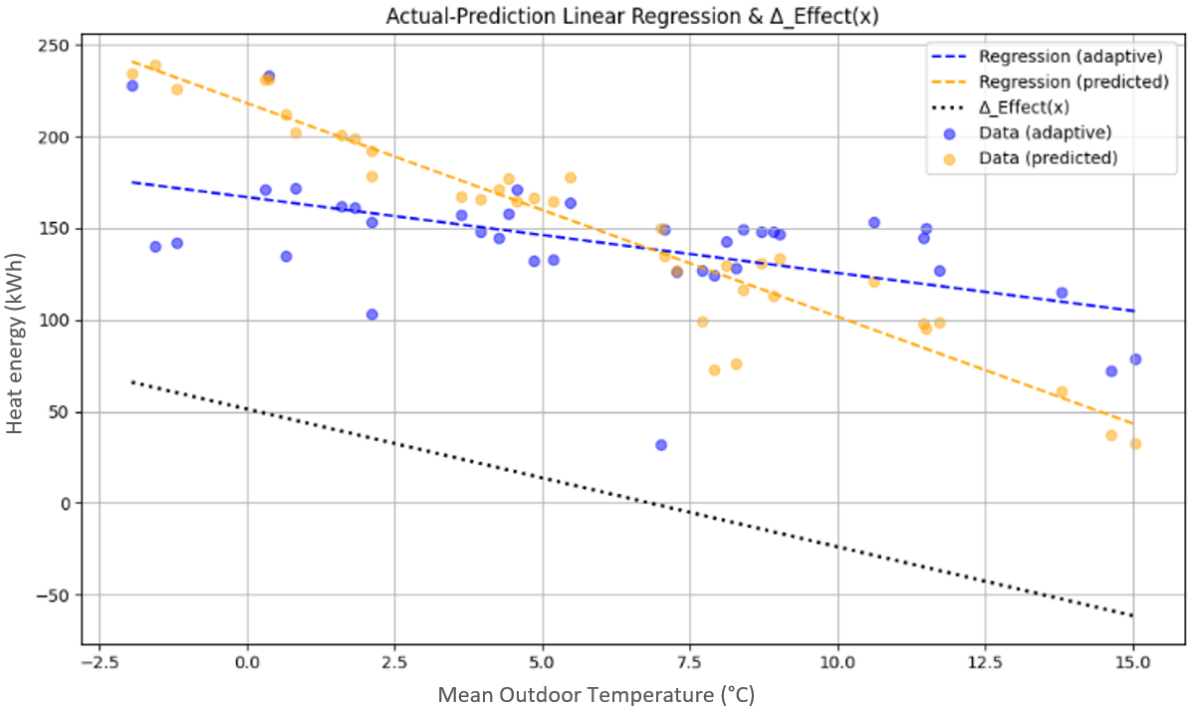

Figure 14: The scatterplot displays the actual data points for the adaptive configuration and the predicted data points for a simulated control group for Bülsweg, including the fitted regression lines: adaptive (blue), simulated (orange) and net effect (black). The first 10 samples from the adaptive configuration were excluded to account for

the initial exploration phase of the optimization. The net effect regression suggests an estimated total energy saving of approximately 40 kWh across the explored temperature range.

3.2.2.4. Digital Twin

Initially, we aimed to use an EnergyPlus-based physical model of the building. However, discrepancies in the heating system's control behavior prevented it from accurately reproducing the real dynamics. The main issue was the reduced heat output near setpoint temperatures. Therefore, we adopted a statistical digital twin approach based on cascaded machine learning models, calibrated and validated on measured data, see figure 15.

A digital twin is considered valid if it

- responds correctly to key inputs: weather, control actions,
- predicts building behavior with sufficient accuracy, and
- remains transparent, reproducible, and documented.

The twin consists of six interconnected models representing successive physical processes, from inlet temperature to indoor comfort. Each model uses weather, control, and hydraulic features to predict its target variable. Validation on unseen data shows consistent predictive performance (R^2 = 0.70 – 0.95 across models) and on the hold-out test set (10.01. - 23.01.2025), the digital twin achieved RMSE = 2.43 kWh, MAE = 1.81 kWh, and R^2 = 0.82, outperforming the GAM baseline (R^2 = 0.69). To assess physical plausibility, we further applied sanity checks using extreme input conditions. Under the assumption of calculated zero inlet temperature, we expect a shutdown of heat transfer throughout the cascade:

$$\text{inlet temperature} = 0 \rightarrow \text{return temperatur} \sim 0 \rightarrow \text{heat energy} \sim 0$$

Within the test window (01.02. – 22.02.2025), the digital twin responded consistently with this expectation, predicting inlet temperatures around 7 °C, return temperatures around 20 °C, and a reduced total heat energy of approximately 1900 kWh. In contrast, baseline regressors collapsed toward their training mean, producing inlet ≈ 35 °C, return ≈ 30 °C, and total energy ≈ 2500 kWh.

A discussion of the full architecture, feature selection process, and detailed performance metrics are provided in Appendix G: Machine Learning-Based Digital Twin and for further details regarding the sanity checks, see H.1.3 Sanity Checks: Perturbation Experiment.

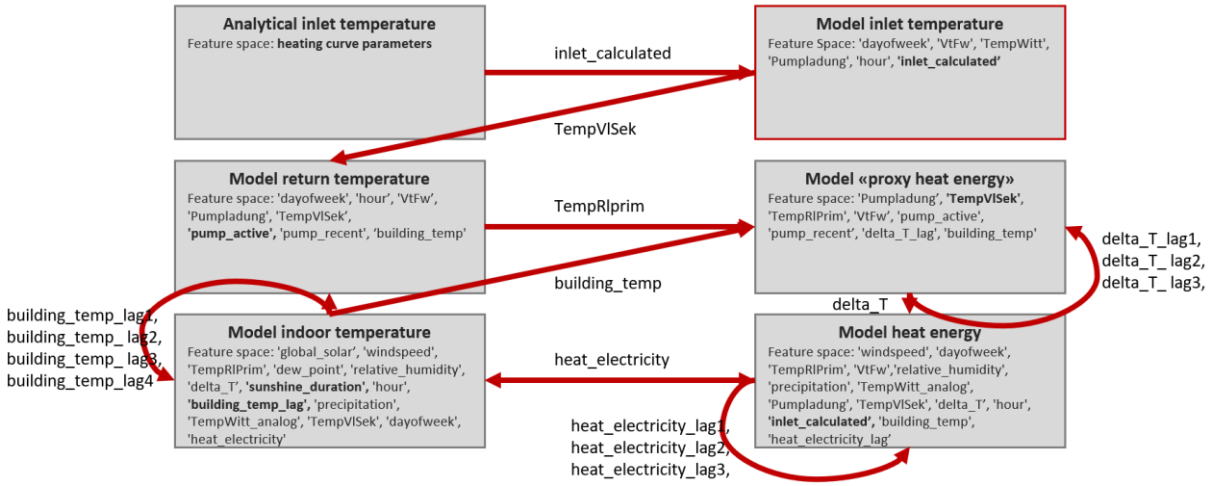

Figure 15: Architecture of cascaded models: The figure illustrates the hierarchical architecture of the data-driven digital twin, composed of six interconnected models predicting hourly temperature and energy dynamics. Each block represents one model, with its primary feature space listed below the block; the most influential input variables are highlighted in bold. The sequence follows the causal flow of the heating process: (1) the Analytical inlet temperature module computes the theoretical supply temperature (inlet_calculated) based on the heating curve parameters; (2) the Model inlet temperature refines this estimate to predict the actual measured supply temperature (TempVlSek), accounting for control noise and pump activity; (3) the Model return temperature predicts the

return water temperature (TempRlPrim), reflecting the thermal energy extracted by the building; (4) the Model proxy heat energy estimates an intermediate proxy variable (ΔT = TempVlSek – TempRlPrim) as a measure of effective heat transfer; (5) the Model heat energy predicts the total thermal energy consumption (heat_electricity); and (6) the Model indoor temperature predicts the average indoor temperature (building_temp), capturing the resulting comfort response. Directed links between the models indicate the flow of predicted variables that serve as inputs for subsequent models, ensuring physical consistency and causal interpretability. For instance, TempVlSek predicted in step (2) feeds into step (3), and TempRlPrim into step (4), while building_temp also influences the proxy energy model to reflect reduced heat transfer at higher indoor temperatures. Lagged variables are used in the proxy, energy, and indoor temperature models to capture temporal dependencies.

Using two digital twins: DT1 (adaptive) and DT2 (static), we estimate the intervention's effect via a bootstrapped Difference-in-Differences approach that accounts for sampling and model uncertainty with respect to our testbed. For a detailed description incl. discussion of accounting for model uncertainty / bias of the digital twin, see Appendix H: Pseudocode of the bootstrapped Difference-in-Differences algorithm.

- **DT1:** Synchronized with the real building and **adapted** with the algorithm after intervention
- **DT2:** Operates without intervention, serving as a control / **static** baseline
- **TB:** measurements of real building Bülsweg
- **pre** (before intervention) and **after** (after intervention).

Accounting for sampling variability: We build a sampling distribution for the difference-in-differences between the two digital twins while injecting model uncertainty and bias estimates. Specifically, DT1 and DT2 are used to predict the hourly heat energy consumption of Bülsweg for the pre (15.11.24 – 22.01.25) and after (02.02. – 17.03.25) periods. These hourly predictions are then aggregated into daily energy usage values, using 5 pm–to–5 pm intervals. The resulting arrays of daily energy values are resampled via a bootstrap procedure to generate a distribution of DID estimates.

$\overline{DT1}_{pre}$ = the mean of daily DT1 predictions in the bootstrap generated array (bootstrap sample)

$\overline{DT2}_{pre}$ = the mean of daily DT2 predictions in the bootstrap generated array (bootstrap sample)

$\overline{DT1}_{after}, \overline{DT2}_{after}$ similarly for after-period

One bootstrap difference-in-difference draw, denoted by (b), is:

$$DID^{(b)} = \left(\overline{DT2}_{after}{}^{(b)} - \overline{DT1}_{after}{}^{(b)}\right) - \left(\overline{DT2}_{pre}{}^{(b)} - \overline{DT1}_{pre}{}^{(b)}\right)$$

And is generated 5000 times. The mean effect is:

$$\overline{DID} = \frac{1}{B}\sum_{b=1}^{B} DID^{(b)}$$

and reported together with the percentile-based confidence interval studied for two cases, Case 1 and Case 2 discussed below.

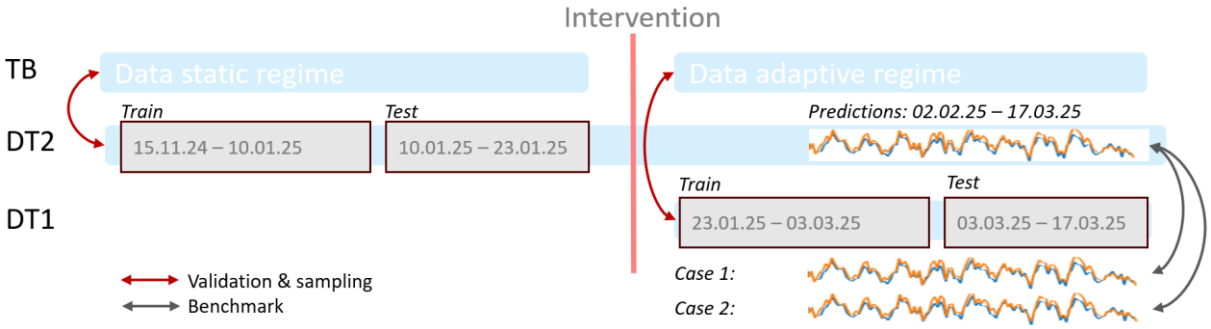

Figure 16: Overview of the DT1–DT2 comparison setup. DT1 and DT2 were trained on distinct operating regimes to represent the adaptive and static heating behaviors of the testbed Bülsweg. DT2 serves as the control digital twin, trained on data collected between November 2024 and January 2025, while DT1 was trained on data from January to March 2025 to reflect the adaptive configuration. A separate hold-out test set was used for validation. Two comparison cases are analyzed: Case 1: DT1 synchronized with the real testbed operation. Case 2: DT1 operating with fully converged Bayesian Optimization parameters and reduced lower bounds (down to 20 °C). The evaluation horizon for both digital twins covers the period 01.02.2025–15.03.2025. The red arrows indicate relationships between a digital twin and the testbed, representing validation (via R^2 and error metrics) and sampling (via mean offset and standard deviation of predictions). The dark grey arrows indicate benchmarking relationships between the two digital twins (DT1 ↔ DT2), comparing their heat energy predictions.

The analysis shows:

- **Case 1** (direct comparison with experiment curve): mean DID = –11.3 kWh/day, 95% CI = (–39.7, +17.5)
- **Case 2** (extrapolated curve lower bound, more flexible dynamic curve): mean DID = –3.7 kWh/day, 95% CI = (–30.9, +24.1)

Interpretation: Statistically, we cannot conclude that our strategy leads to significant energy savings, as the confidence interval overlaps with zero.

In Case 1, we used DT2 to predict the potential daily heat energy consumption for the period 02.02. – 17.03.25, i.e., after the calibration phase of our algorithm and in synchronization with the applied inlet temperatures for the real building (see Section 3.2.6.1 Results Bülsweg: Convergence & Robustness). Considering the known limitations discussed in Section 3.2.6 Discussion Energy Net Savings and illustrated in Figure 17, the model indicates an average loss of 11 kWh/day, with possible outcomes ranging from a loss of 40 kWh/day (worst case) to a gain of 18 kWh/day (best case).

In Case 2, we compared DT2 and DT1, where DT1 was reparametrized to match the static heating curve configuration (see table in Section 3.2.6 Discussion Energy Net Savings). The parameters $(\mathbf{X_3}; \mathbf{Y_3})$ and $(\mathbf{X_4}; \mathbf{Y_4})$ were adjusted such that the curve could reach temperatures as low as 20 °C, depending on outdoor conditions. Furthermore, the converged heating tuner predictions were used to estimate the optimal $\mathbf{Y_2}$. Under this setup, the results show near parity between the static and adaptive configurations, indicating that the new parameterization effectively mitigates most of the previous energy penalty.

Appendix I provides the full pseudocode for evaluating heat energy savings using DT1 and DT2, including the implemented uncertainty treatment.

3.2.2.5. Results Bülsweg: Convergence & Robustness

We perform an empirical convergence analysis in the setting where the true objective function is unknown. Specifically, we use a local convergence test that evaluates the stability of the model's predictions within discretized context–action squares. The analysis is based on GPR models trained on incrementally larger datasets.

Experimental Setup

- We define a grid consisting of 54 context – action squares, determined by:

- Context values: [-10, -7.5, -5, -2.5, 0, 2.5, 5, 7.5, 10, 12.5]
- Action values: [31, 34, 37, 40, 43, 46]

- For each square, we train a new model with progressively increasing data sizes from 10 till 60 by additional 5 datapoints: [10, 15, 20, 25, 30, 35, 40, 45, 50, 55, 60].
- At each stage, we predict the mean and standard deviation at the center point of each square and observe how predictions evolve with more data.

Convergence after 40 datapoints

At low data densities (e.g., 10 data points for 56 squares → ~0.18 points/square), the model's prediction surface is relatively flat and only sensitive in areas with lower actions. As more data is added, particularly in underrepresented regions (e.g., context 0, action 46), significant shifts in predicted means occur.

We visualize this progression by inspecting prediction grids over time. The shift in the prediction values becomes apparent when comparing early stages (with almost constant combined score values) to later stages (where the combined scores exhibits structure and local extrema).

We observe that after approximately 30 – 40 datapoints for our setting, the predictions stabilizes for the already explored areas and there is no significant improvement, suggesting that the model has converged to a robust heating policy.

To understand the model's exploration behavior over time, we visualize the frequency of predictions across a discretized 6×9 context – action grid. The bar on the right represents how often each square in the grid was selected by the model, the values inside each square are the predictions for the mean combined score and it's standard deviation at three stages: early (10 datapoints), intermediate (40 datapoints), and late (60 datapoints). For example, the square at Action Bin 37 and Context Bin 0.0 contains 3 values, with a predicted mean combined score of 1.19 and a standard deviation of ±0.04. Even when standard deviation estimates are low in early phases, the mean prediction may still vary significantly, suggesting that small uncertainty does not guarantee convergence in low-data regimes.

Early Stage (10 datapoints): uniform values across squares.

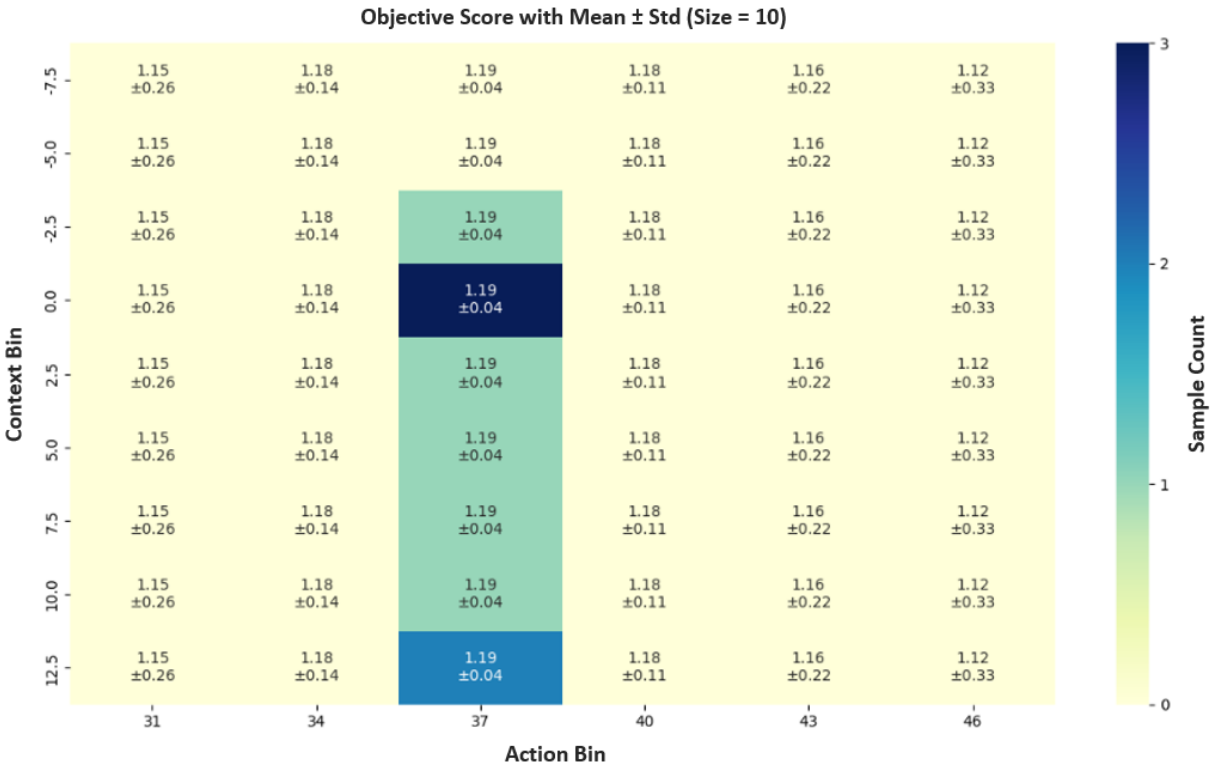

Middle Stage (40 datapoints): emerging variation, e.g., square at (46, 0) shows a notable change.

Final Stage (60 datapoints): smoother, more structured values indicating a more refined model.

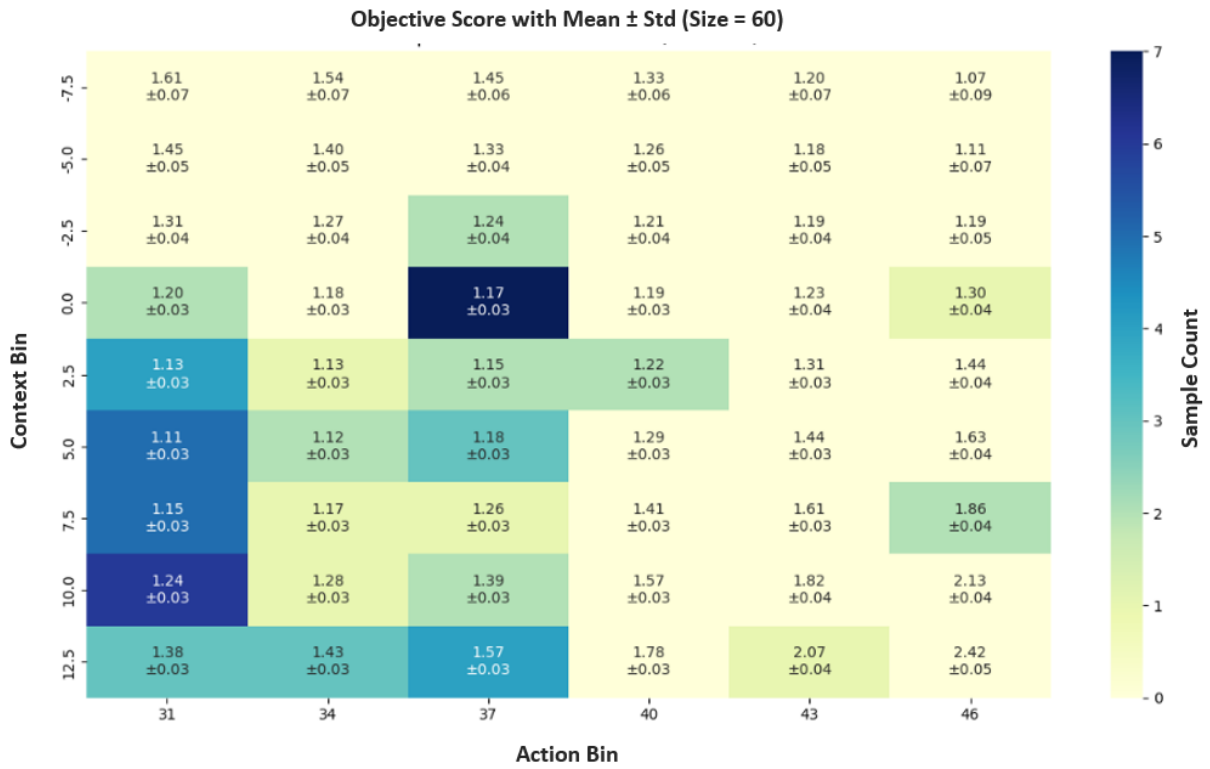

The increasing density and spread of samples across the grid demonstrate progressive exploration followed by exploitation. Notably, high-utility regions (e.g., column 2 across several rows) are consistently reinforced, while seemingly low-performing regions remain unvisited. This transition is controlled by adjusting the relative weights of exploration and exploitation of our BO implementation. As more data points are added, the relative importance of exploitation increases, while exploration decreases, speeding up

convergence. This dynamic weighting is a distinctive feature of our implementation, differing from standard BO.

To quantify convergence, we compute the average prediction change across all squares of this grid between consecutive training steps. The results can be looked up in figure 20 and the visualization reveals:

- High instability in early iterations (especially step 1 and 2, from 10 to 15 and 15 to 20 data points).
- A stark reduction in mean prediction variation after 30 – 40 datapoints, indicating convergence behavior.
- Even when standard deviation estimates are low in early phases, the mean prediction may still vary significantly, suggesting that small uncertainty does not guarantee convergence in low-data regimes.

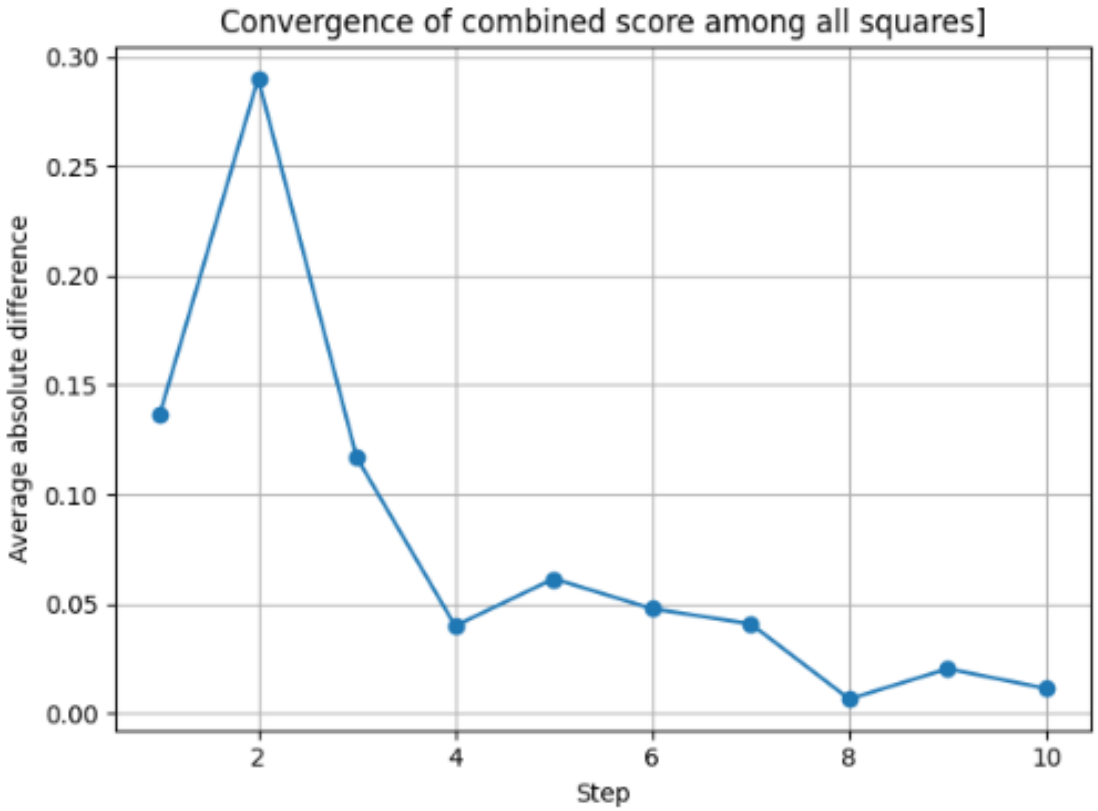

Figure 20: Plotted are the average differences between the predictions made for different steps: Step 1 stands for the average difference between 10 to 15 datapoints, step 2 stands for the average difference between 15 to 20 datapoints, step 3 stands for the difference between 20 to 25 and so forth. The average prediction change reveal: high instability in early iterations (especially from 10 to 15, and 15 to 20 data points). A stark reduction in mean prediction variation after 30 – 40 datapoints, indicating convergence behavior.

Robustness ok for positive outliers bad for negative outliers

To evaluate the robustness of our model when optimizing the heating curve, we manipulated existing samples and added them to the data used to train our model. Specifically, we randomly selected 8 distinct indices from our original dataset, namely: [3, 15, 24, 31, 37, 41, 50, 57] and multiplied their objective score by a distortion factor. The specified samples were distorted by a factor of 3 and replaced the old value. A factor of 3 corresponds approximately to a standard deviation of 3 with respect to our assumed normally distributed data, which classifies them as strong positive outliers. After each replacement, the model was retrained, and its prediction was evaluated across a fixed set of contexts ([-10, -7.5, -5, -2.5, 0, 2.5, 5, 7.5, 10] °C) and actions (inlet temperatures ranging from 28 °C to 45 °C in 0.1 °C

steps), see figure 21. We did the same for a factor of 0.1 to obtain negative outliers and mixed between factor of 3 and 0.1.

The tables below summarize the model's predicted optimal inlet temperature per context as the number of manipulated samples increases (Case 0: no distortion, Case 6: 6 distortions added):

Factor 3:

Context (°C)	Case 0	Case 1	Case 2	Case 3	Case 4	Case 5	Case 6
-10	45	45	45	43.11	43.11	45	45
-5.0	43.11	45	39.33	39.33	39.33	39.33	41.22
-2.5	37.44	41.22	35.56	35.56	35.56	37.44	39.33
0.0	33.67	33.67	33.67	33.67	33.67	35.56	37.44
2.5	29.89	28	31.78	31.78	31.78	33.67	35.56
5.0	28	28	29.89	29.89	31.78	31.78	31.78
10	28	28	28	28	29.89	29.89	28

Factor 0.1 and Mixed:

Context (°C)	Case 0	Case 1, 0.1	Case 4, 0.1	Case 6, 0.1	Case 6, Mix
-10	45	45	45	45	45
-5.0	43.11	45	45	28	45
-2.5	37.44	35.56	45	28	45
0.0	33.67	31.78	45	28	45
2.5	29.89	31.78	45	28	45
5.0	28	31.78	45	28	45
10	28	31.78	29.89	45	28

The results reveal a critical asymmetry in the system's robustness:

- Factor 3: seem to add noise, but they do not remove the signal, the model could still reliably converge toward a reasonable solution.
- Factor 0.1: seem to destroy the signal, if we reduce the rewards to near zero for some samples, the model predictions flatten, the model believes that almost no actions yield to any optimal solution. After enough negative distortions have been added the model starts to make bad decisions. Concrete example for a bad decision is observed in Case 6 (factor 0.1) where at 10°C, the system wrongly predicts the maximum action of 45°C when it should recommend much lower inlet temperatures based on the undistorted case of 28°C.
- Mixed: seem that the negative effects dominated. Although half of the distortions were positive, the system still suffered from bad decisions, reinforcing the finding that underestimations are more harmful than overestimations.

Case 1:

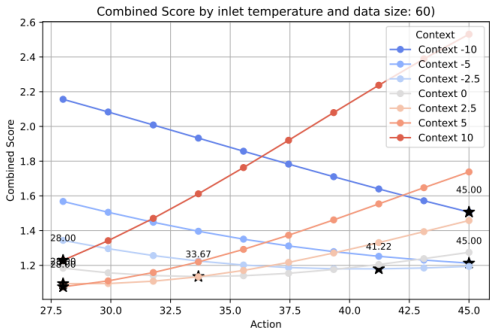

Case 2:

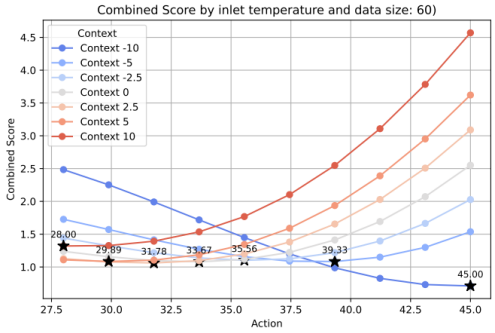

Case 3:

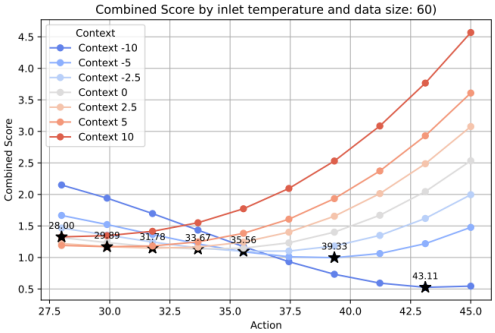

Case 4:

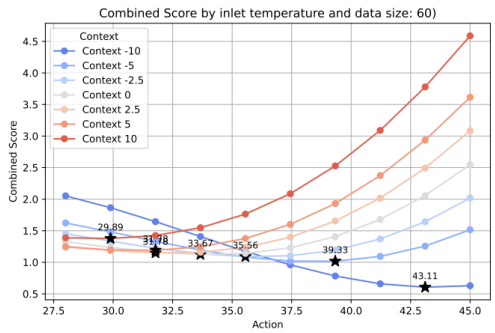

Case 5:

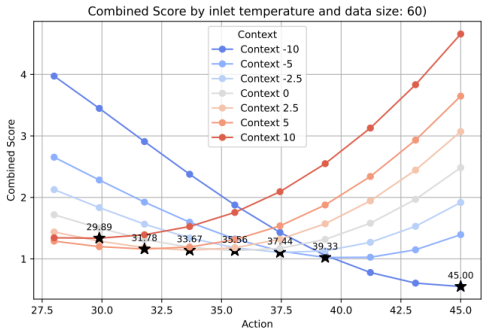

Case 6:

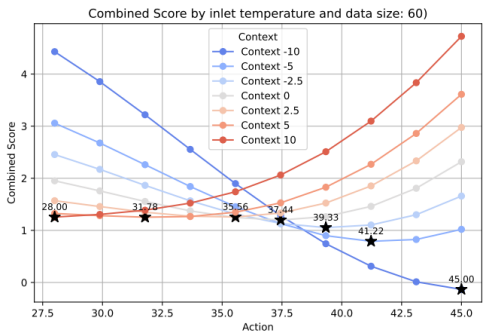

Case 0:

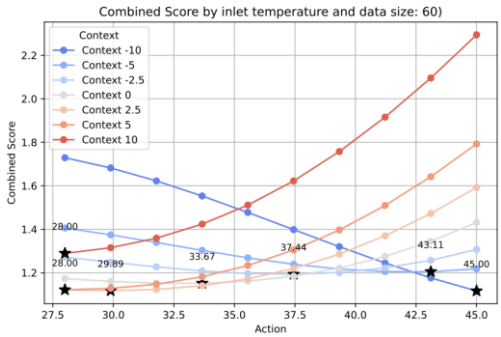

Figure 21: Model predictions for fixed contexts across the full range of inlet temperatures (actions) are shown for varying numbers of distorted data points, from case 0 to case 6. Notably, the curve progression in case 1 stands out from the others, which exhibit more consistent and similar behavior.

3.2.2.6. Discussion Energy Net Savings

Method	Energy Net Savings
Linear Regression	By comparing the average energy consumption before and after the intervention and fitting a linear regression to sampled data points, we estimated a **net saving of 5.7%**.
HDD	Using the HDD normalization technique, which adjusts for weather variability, we calculated a **net saving of 4.1%.** This method is robust for such comparisons and likely provides the most conservative and reliable estimate.
Machine Learning	We trained an ML model to predict energy consumption based on weather and room conditions, using it as a virtual control group. The model achieved an average **net saving of 1.6%.** This approach to estimate
Digital Twin DID	**No statistically significant effect** was observed (confidence intervals: –39.7 to +17.5 and **–30.9 to +24.1**). Based on the manual alignment with the static heating curve in **Case 2**, we had expected comparable performance under mild temperature conditions and reduced energy use during colder periods. However, with the current model, this hypothesis cannot be confirmed.

While linear regression, HDD and Machine Learning point to a positive impact of the algorithm, the HDD-based estimate is likely the most reliable due to its ability to normalize for weather influences without being overly sensitive to model assumptions. The ML model supports this finding, validating the presence of a small but measurable effect.

The most plausible explanation for the discrepancy between the Digital Twin results and the other evaluation methods lies in the training characteristics and extrapolation limitations of the machine learning models used to construct the digital twin. These models are optimized to reproduce observed system behavior within the range of their training data, meaning their boundaries and response functions are tuned to minimize error for known patterns. As a result, they perform well within the data domain but struggle to extrapolate to unseen parameter configurations, such as heating curve settings that were not represented during training (see Figure 18).

In our case, DT1 was trained predominantly on data corresponding to a heating curve parameterization that did not include values below 20 °C, which reflects the system's historical operating range. Consequently, DT1 lacks exposure to configurations representing milder outdoor conditions and lower inlet temperatures, making it difficult to accurately predict potential efficiency gains in those regimes.

Additionally, the Bayesian Optimizer was tuned within a restricted temperature range (30 – 45 °C), further limiting exploration of potentially more energy-efficient configurations below 30 °C. Consequently, while real-world measurements confirm that the adaptive strategy yields modest energy savings (~5 %), the current statistical digital twin framework cannot yet reproduce this outcome consistently due to model asymmetry and extrapolation bias (DT1, BO). This asymmetry, manifested as unequal sensitivity to lower versus higher inlet temperatures and corresponding changes in heating energy, is further discussed in Section H.1.3, Sanity Checks: Perturbation Experiment.

We believe that the actual saving potential is positive and probably higher than the estimated savings estimated for the HDD and Linear Regression method. This is mainly because our current heating curve configuration imposes a lower limit of 30°C for the inlet temperature, even during periods of mild outdoor conditions. In contrast, the previous static heating curve configuration allowed the inlet temperature to drop to as low as 20°C, resulting in lower energy consumption.

We observe inefficiencies in the high temperature range. These inefficiencies are believed to be caused due to the configuration of our heating curve in comparison with the old configuration. Our configuration didn't allow to go below 30 degrees, while the old configuration allowed to go down till 20 degrees, see table below. The table below highlights this difference in configuration.

Point	Old Setting (Temperature °C)	Our Configuration (Temperature °C)
$(\mathbf{X_1}; \mathbf{Y_1})$	(-10; 36)	(-20; 36)
$(\mathbf{X_2}; \mathbf{Y_2})$	(0; 32)	(-10; **variable**)
$(\mathbf{X_3}; \mathbf{Y_3})$	(10; 28)	(10; 30)
$(\mathbf{X_4}; \mathbf{Y_4})$	(20; 20)	(20; 30)

A rough estimation of incurred energy losses can be made by comparing the theoretical inlet temperatures from the adaptive and static heating curves under mild conditions, as shown in Figure 12. We looked at a warm day with an average outdoor temperature of around 13 °C. The area under the inlet temperature curves, which is proportional to heating energy, suggests a loss of approximately 16% with respect to the adaptive configuration, see figure 17.

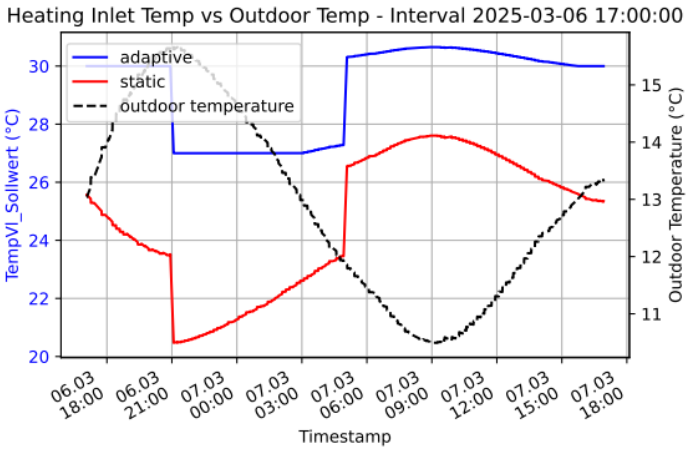

Figure 17: We plotted the theoretical inlet temperatures based on outdoor temperature and heating curve parameterizations. While the adaptive curve (TempVl_Sollwert) can be directly retrieved from the heating system, the inlet temperatures for the static configuration were computed for comparison, red curve. The difference in the area under both curves, proportional to heating energy, amounts to approximately 16% (= $\frac{area_{adaptive}}{area_{static}}$).

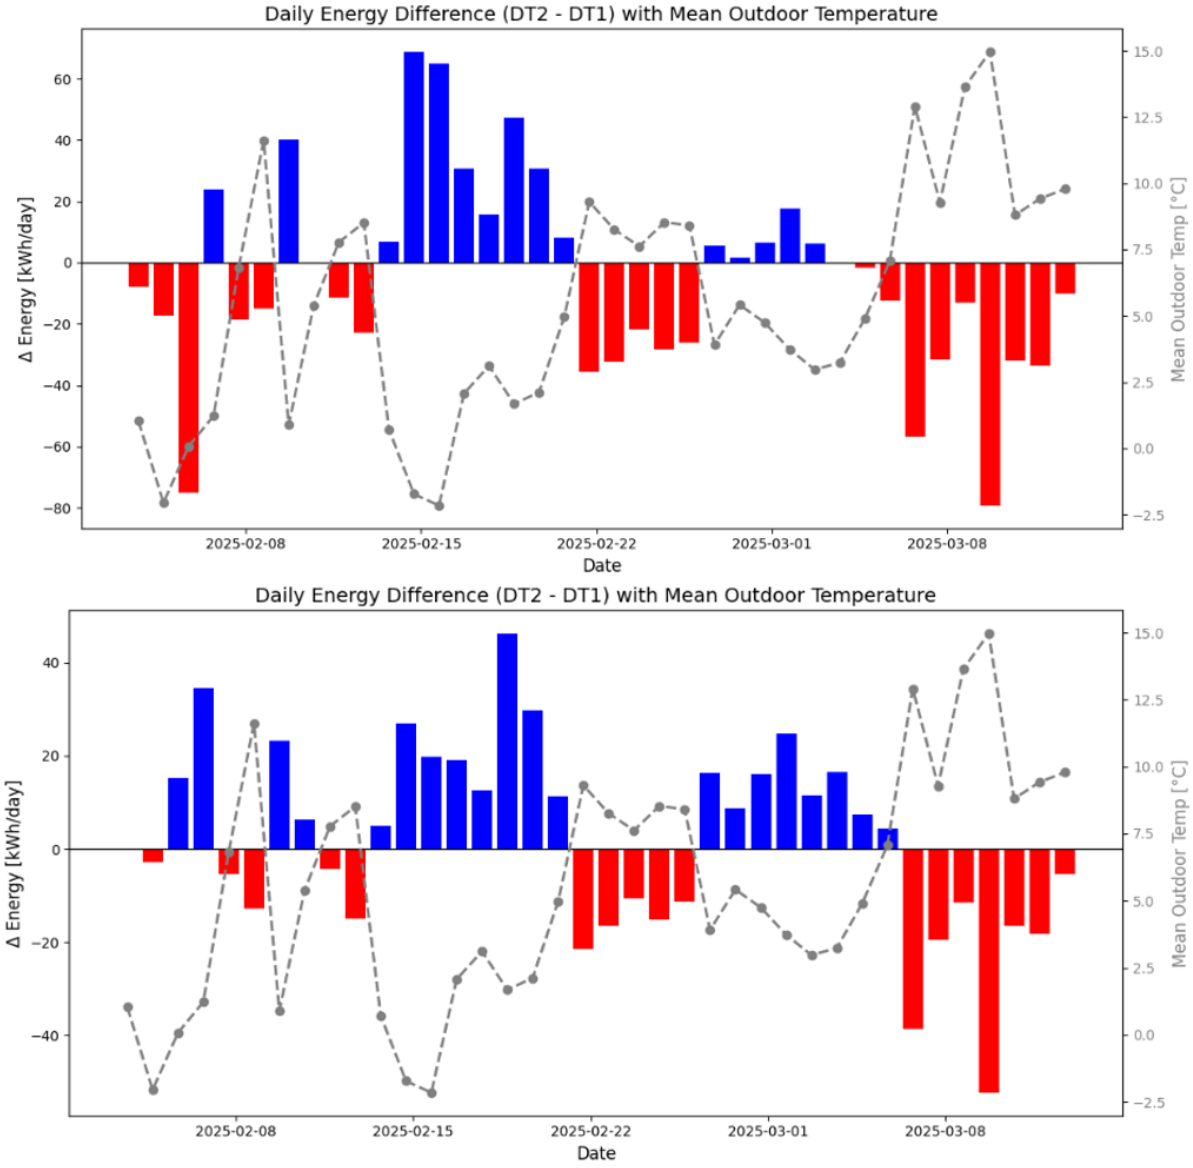

Figure 18: Daily differences in predicted heat energy consumption between the static (DT2) and adaptive (DT1) digital twins. Positive (blue) bars indicate energy savings of the adaptive strategy, while negative (red) bars represent increased energy use relative to DT1. The gray dashed line shows the mean outdoor temperature for the corresponding days. A consistent pattern emerges across both scenarios: at mild outdoor temperatures, the adaptive control tends to consume more energy than the static configuration, while at lower temperatures, it achieves higher efficiency and energy savings. The upper panel corresponds to **Case 1**, where DT1 is synchronized with the real testbed, and the lower panel to **Case 2**, where the lower bound of the heating curve was reduced to 20 °C and the predictions were derived from a fully converged Bayesian Optimization configuration.

Comfort: static comfort score 0.24 vs. adaptive comfort score of 0.23

We now turn our focus to thermal comfort. To assess it, we estimated the daily comfort score for both the static and adaptive heating configurations using equation (1.1). This calculation was based on indoor temperature measurements from the living spaces in each apartment. Due to irregular and sometimes missing data, we included only those apartments with at least one valid reading during both day and night periods. For each qualifying sensor reading, we calculated the deviation from the comfort thresholds: 23 °C during the day and 20 °C at night. These deviations were averaged per apartment and then aggregated across all apartments to estimate the overall building-level comfort.

The resulting comfort score values were binned for visualization, see figure 19, and we fitted a Gaussian distribution to the histograms for both heating configurations. Based on the average comfort score, we find that comfort levels remained nearly unchanged after switching from the static to the adaptive heating curve, demonstrating that thermal comfort for residents is maintained.

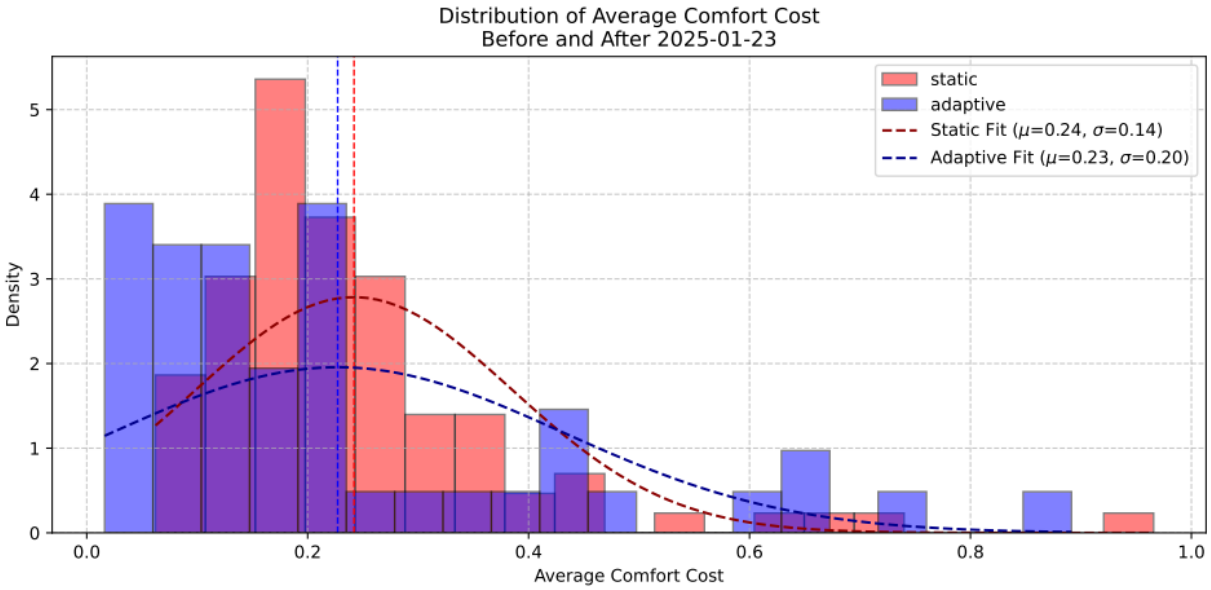

Figure 19: The histograms show the distribution of daily comfort score per apartment for both heating configurations. A Gaussian distribution is fitted to each histogram to highlight the mean and spread of comfort scores. The comparison shows that the average comfort score remains nearly unchanged between the static and adaptive configurations, indicating that occupant thermal comfort is maintained following the implementation of the adaptive heating curve.

3.2.3.Field Test NEST Sprint

3.2.3.1. Building Description and Setup

The second field deployment was conducted at NEST Sprint, a 167 m² office unit at the NEST research platform at Empa Dübendorf. The unit comprises 12 individual offices and was opened in August 2021, built predominantly from reused materials and components. Heating energy is sourced from the NEST Medium Temperature Energy (MTE) grid via a heat exchanger and distributed via a ceiling heating/cooling system. Room temperatures are individually controllable via one valve per room. The supply temperature is governed by a weather-compensated flow temperature control: the setpoint is calculated as a function of outdoor temperature via an adjustable heating curve — the parameter optimized by AHA. The unit faces primarily east, with limited south-facing window area (~7 m² total), restricting direct passive solar gains to morning hours. The heating system is accessible via the NEST API and MQTT interface, see Appendix I.

The adaptive heating curve was activated on February 24, 2026, following a static reference period from October 1 to November 29, 2025 (60 days, Yb = 32°C constant). Run_003 produced 34 daily samples in total: 5 initialization days (February 24–28, Yb = 31.0, 32.25, 33.5, 34.75, 36.0°C) and 29 autonomous BO decisions (March 1–30, 2026).

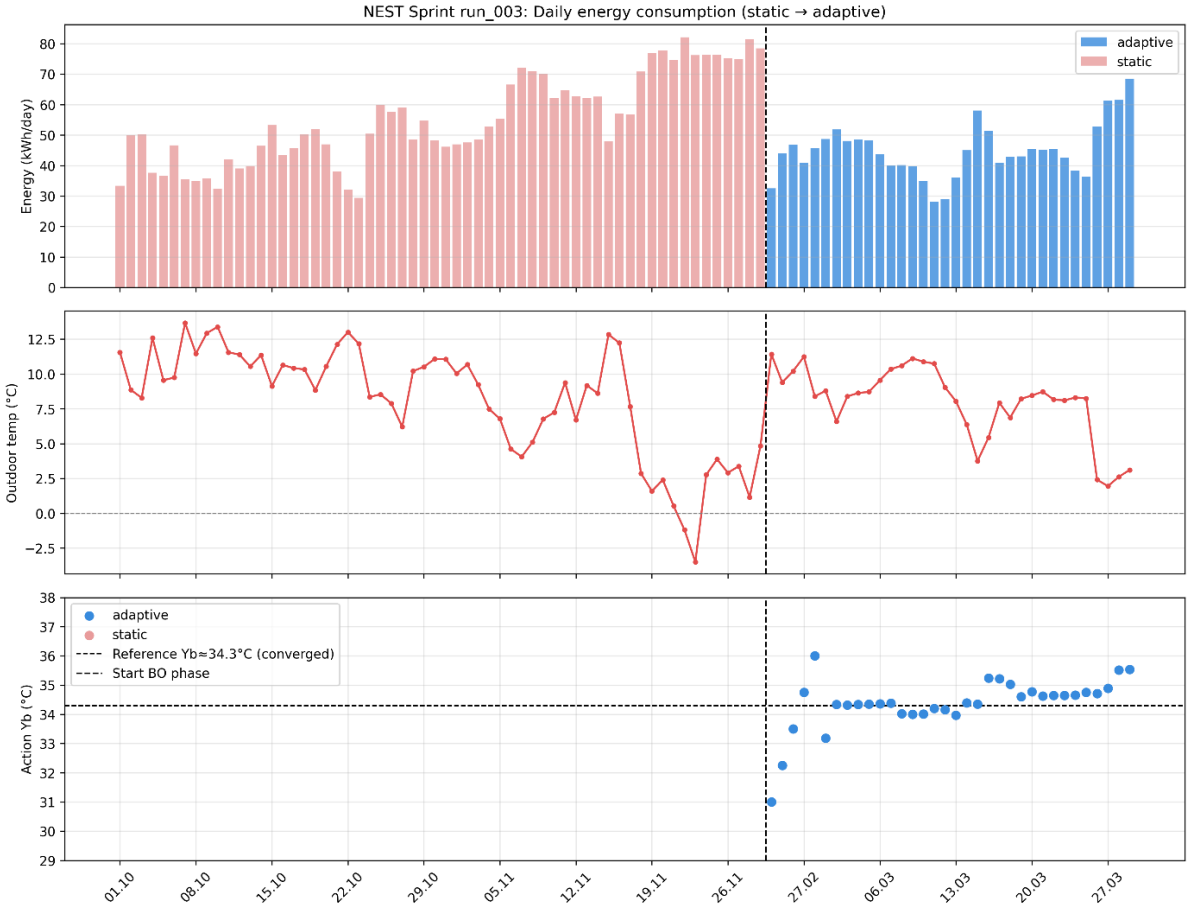

Figure 20: overview of the full measurement campaign (October 2025 – March 2026). Top: daily heat energy (red: static, blue: adaptive). Middle: outdoor temperature. Bottom: Yb setpoint with converged reference at 34.3°C. Dashed vertical line marks BO phase start on March 1, 2026.

3.2.3.2. Benchmark Analysis

The evaluation strategy for NEST Sprint differs from the Bülsweg deployment (Section 3.2.2.1). In Bülsweg, the static reference and the adaptive phase covered consecutive segments of the same heating season, allowing a direct comparison without seasonal confounding. In Sprint, by contrast, the static reference period (October–November 2025) and the adaptive phase (February–March 2026) fall in different seasons. Mean solar irradiance in February–March is approximately double that of October–November, and additional structural differences for example daylight duration, introduce a confounding effect that simple model-based extrapolation cannot reliably separate from the AHA contribution. To address this, we adopt a stepwise subtractive approach: we first estimate the raw energy savings between the two phases, then progressively remove the contribution of known confounders such as solar irradiance and seasonal background variability. The residual difference between the static and adaptive phase, after these corrections, is interpreted as the AHA-attributable effect. The evaluation proceeds in three methodologically independent steps:

(1) Regression analysis (Section 3.2.3.3, methodology following Section 3.2.2.1) — fits separate linear models of daily heat energy on outdoor temperature for the static and adaptive phases. Yields the headline saving of 17.6% (raw mean difference) at the observed temperature distribution.

(2) Heating Degree Day analysis (Section 3.2.3.4, methodology following Section 3.2.2.2) — normalises consumption by HDD to control for differences in outdoor temperature. Yields 22.4% ± 0.7% under range-restricted bootstrap.

(3) Control period comparison (Section 3.2.3.5) — introduces two historical February–March periods (2024 and 2025) without AHA as seasonal baseline. Both periods were verified as pure static operation – no experimental activity over 2752 hourly records. This step quantifies the seasonal baseline shift directly from observed data, without relying on model assumptions, and produces the final AHA net effect: 5 – 10 %.

3.2.3.3. Regression Analysis

The regression methodology follows Section 3.2.2.1 (Equations 3.0–3.2). Unadjusted comparison: mean daily energy consumption decreases from 55.1 kWh/day (static, n=60) to 45.4 kWh/day (adaptive, n=29) — a raw reduction of 17.6%.Temperature-adjusted regression:

$$y_{static} = -3.37 * x + 82.5 \ (r^2 = 0.81)$$
$$y_{adaptive} = -2.90 * x + 67.4 \ (r^2 = 0.66)$$
$$\Delta_{slope} = -0.47$$
$$\Delta_b = 15.1$$

For example, at $\bar{T}$ = 8.0°C, the regression-predicted saving is 20.3% (ΔE = 11.4 kWh/day). A notable feature is the significantly flatter slope of the adaptive phase (−2.90 kWh/°C) compared to both the static period (−3.37) and the historical control periods (−4.64 and −4.66), potentially indicating that AHA disproportionately reduces energy consumption on cold days, see figure 21 and figure 24.

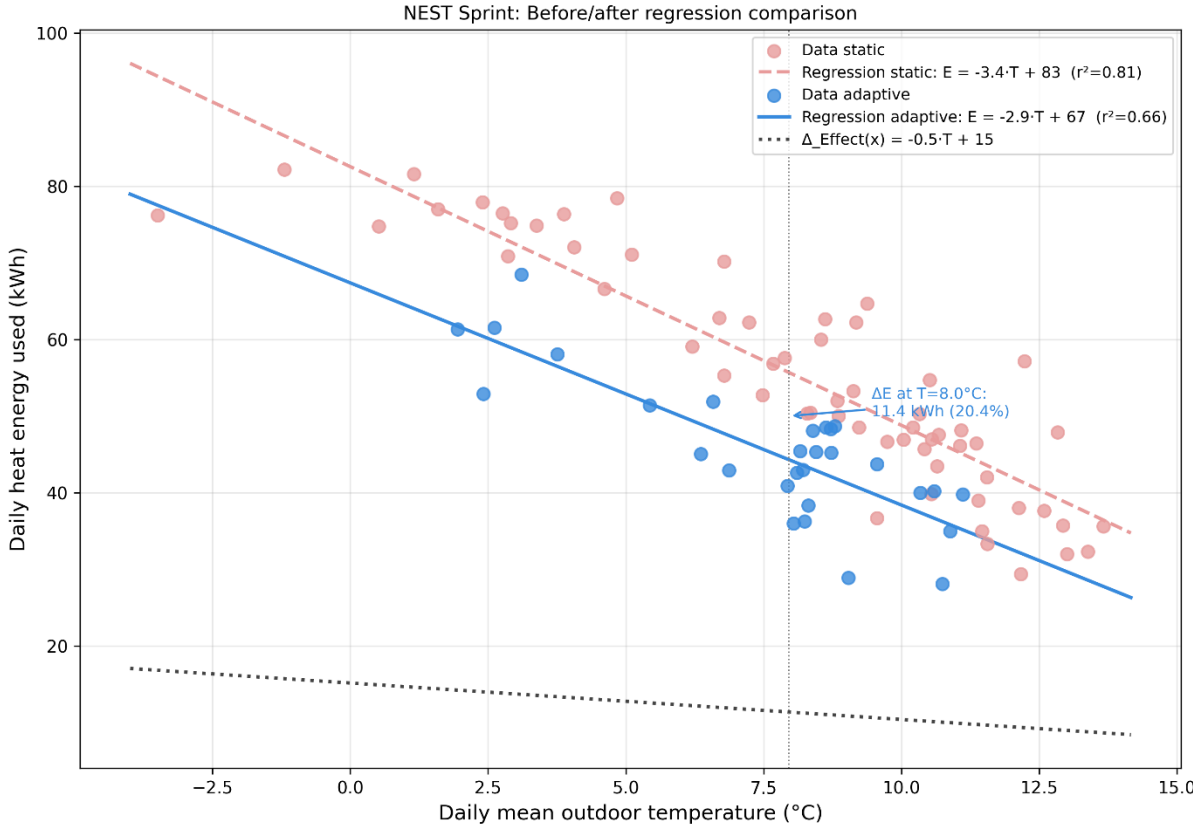

Figure 11: The scatterplot displays the analyzed data points for NEST Sprint, including the fitted regression lines for the static (red), adaptive (blue), and net effect (grey) configurations. To ensure a fair comparison and minimize extrapolation errors, we only included data points from the pre- and post-intervention periods that fall within the temperature range where the adaptive heating curve was active. Additionally, the first 5 samples from the adaptive configuration were excluded to account for the initial exploration phase of the optimization. The net effect regression suggests an estimated total energy saving of approximately 132 kWh across the explored temperature range, corresponding to about 12 kWh per degree Celsius.

3.2.3.4. Heating Degree Day Method

HDD normalization follows Section 3.2.2.2 with T_base = 22.0°C. Specific energy consumption decreases from 3.97 kWh/HDD (static) to 3.20 kWh/HDD (adaptive). Bootstrap-estimated saving (1,000 iterations): -22.4% ± 0.7%. However, this estimate conflates the AHA effect with a structural seasonal difference. The control period comparison below isolates this component.

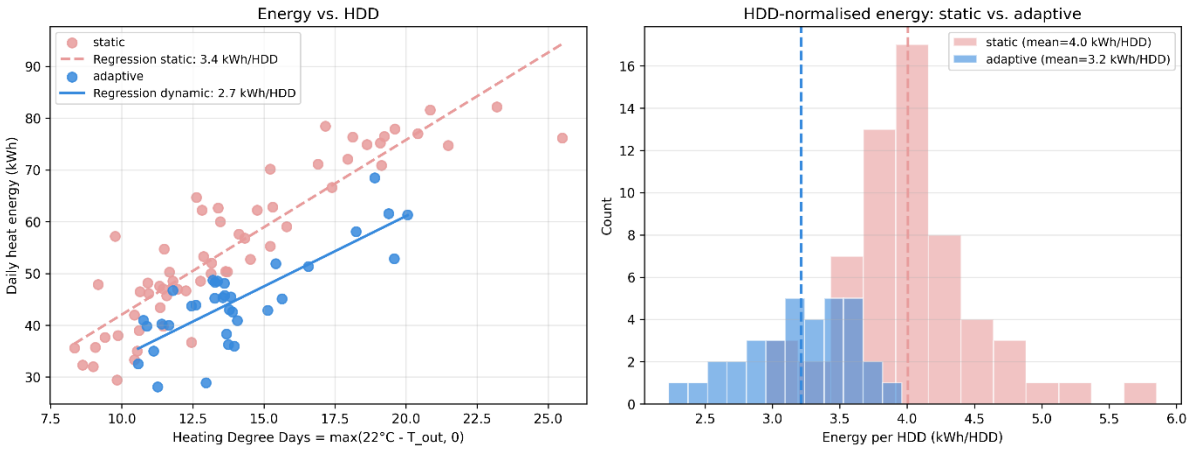

Figure 23: HDD analysis. Left: daily heat energy vs. HDD (base T = 22°C), static and adaptive with regression lines (3.4 and 2.7 kWh/HDD). Right: histogram of HDD-normalized energy intensity (static mean 4.0, adaptive mean 3.2 kWh/HDD).

3.2.3.5. Control Period Comparison

Two historical control periods: February–March 2024 (n=57) and February–March 2025 (n=57), serve as seasonal baselines. Both were verified as pure static operation without running experiments. The solar confounder is substantial: irradiance in the adaptive phase (4,900 ± 1,632 Wh/m²/day) is approximately double the static reference (2,343 ± 1,322 Wh/m²/day). The average indoor temperature was practically identical across static (22.48°C) and adaptive (22.19°C) periods, confirming that solar irradiance acts primarily through facade warming rather than direct room heating.

For each of the four periods (static, control 2024, control 2025, adaptive), we fit a univariate linear regression of daily heat energy on daily mean outdoor temperature "$E(T) = \alpha * T + \beta$", see figure 24.

The daily mean outdoor temperature of the adaptive phase, T_ref = 7.6°C, serves as a common evaluation point. Predicted energy demand at this temperature for each period:

Period	E(T_ref = 7.6°C)	Δ vs. Static
Static (Oct/Nov 2025)	56.9 kWh/day	—
Control Feb/March 2024	51.3 kWh/day	−10.0%
Control Feb/March 2025	49.8 kWh/day	−12.5%
Adaptive (Feb/March 2026)	45.4 kWh/day	−20.3%

The two control periods reveal a structural seasonal baseline shift of 10–12% between autumn and spring due to: longer daylight hours, increased solar irradiance.

Subtracting the seasonal baseline from the observed adaptive saving yields the net AHA effect:

- vs. Control 2024: 20.3% − 10.0% = 10.3%
- vs. Control 2025: 20.3% − 12.5% = 7.8%

To further verify and tighten this estimate, we fit a multivariate linear regression on the static reference period using outdoor temperature, global solar radiation and wind speed as predictors:

$$\mathrm{E} = 86.28 - 3.20 \cdot \mathrm{T} - 0.0013 \cdot \mathrm{Solar} - 2.01 \cdot \mathrm{Wind}\ (\mathrm{R}^2{=}0.83)$$ mit: T in **°C,** Solar in Wh/m²/Tag und Wind in m/s

This model is then applied to each daily record of all four periods to predict the energy that the static-mode building would have consumed under the observed weather. The aggregate residual per period — defined as $(E_{\mathrm{obs}} - E_{\mathrm{pred}})/E_{\mathrm{pred}}$— quantifies the deviation of actual consumption from the static-mode expectation under that period's weather conditions:

Period	E_obs (mean)	E_pred (mean)	Residual
Static	55.1 kWh	55.1 kWh	—
Control 2024	47.5 kWh	53.0 kWh	−10.4%
Control 2025	58.4 kWh	62.1 kWh	−6.0%
Adaptive	45.4 kWh	53.7 kWh	−15.5%

The non-zero residuals of the control periods (−10.4% and −6.0%) reveal a systematic seasonal model bias: the static-trained regression structurally overestimates spring heating demand even in the absence of any AHA intervention. This bias is not a flaw of the model, but an inherent limitation of extrapolating an autumn-fitted relationship to spring conditions (see discussion 3.2.2.4 Digital Twin).

The AHA net effect is the residual of the adaptive phase minus the seasonal model bias measured on the controls:

- vs. Control 2024: −15.5% − (−10.4%) = −5.1%
- vs. Control 2025: −15.5% − (−6.0%) = −9.5%

Convergence and reported interval
Both methods, using independent assumptions, yield consistent results:

Method	Vs. Control 2024	Vs. Control 2025
Direct comparison at T_ref	10.3%	7.8%
Multivariate cross-validation	5.1%	9.5%

Across both methods and both control periods, the AHA-attributable net saving converges to a range of 5–10%.

The remaining uncertainty cannot be reduced further within the available data. A future study with a continuous deployment spanning a full heating season would tighten this estimate considerably.

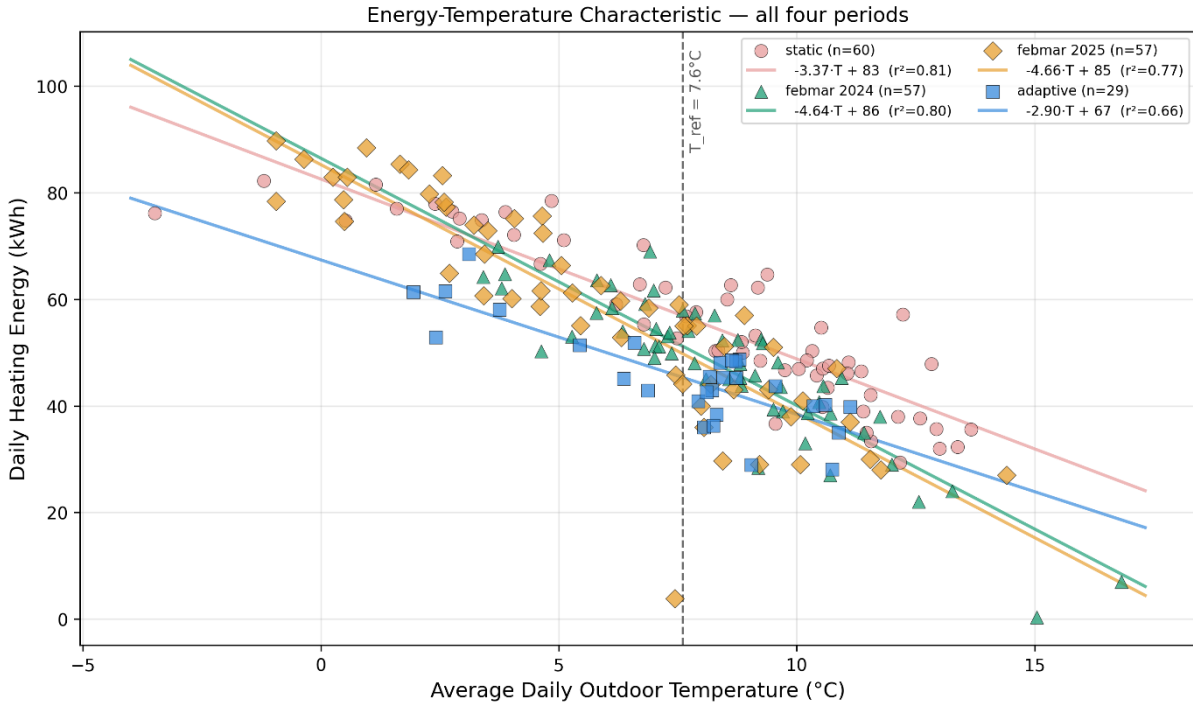

Figure 24: Energy–temperature regression for all four periods. Static Oct–Nov 2025 (red circles), Control Feb–Mar 2024 (green triangles), Control Feb–Mar 2025 (orange diamonds), Dynamic AHA Feb–Mar 2026 (blue

squares). Static lies above the seasonal controls (10–12% baseline shift); AHA adaptive lies below both controls.

3.2.3.6. Convergence and Robustness

Context–Action correlation: over 29 autonomous BO decisions, the Pearson correlation between daily mean outdoor temperature and chosen Yb setpoint is r = −0.537 (p = 0.003): colder days consistently received higher supply temperature setpoints, confirming a physically correct context-dependent heating policy.

Convergence behavior under similar contexts: beyond the formal convergence assessment, additional evidence for stable learning emerges from the BO's response to recurring outdoor conditions. Towards the end of the deployment (from approximately March 20 onwards, see Figure 20, bottom panel), the algorithm consistently produced very similar Yb setpoints (~34°C) for similar outdoor temperatures, indicating that the GP surrogate had reached a stable representation of the local optimum within the operating range. Only when the outdoor temperature dropped to approximately 2.5°C did the algorithm respond with a clearly elevated setpoint (~35.5°C) — a context-dependent reaction that is both physically correct and consistent with the negative slope of the context–action relationship reported above. This pattern is qualitatively comparable to the behavior observed at Bülsweg (Figure 11) and provides evidence for stable convergence even at the lower sample size available at Sprint.

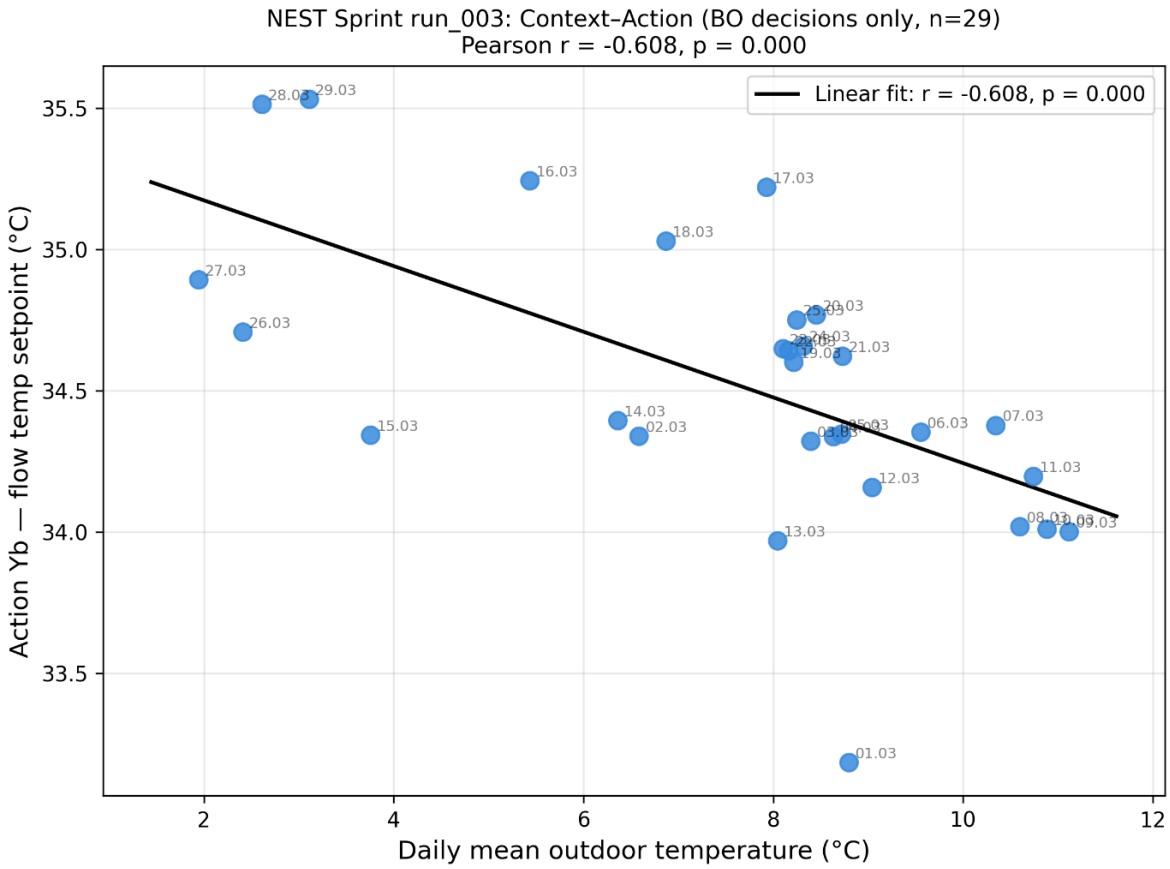

Figure 25: Context–Action scatter plot, BO phase (n=29). Outdoor temperature (x) vs. Yb setpoint (y). Linear fit: r = −0.537, p = 0.003. Colder days consistently receive higher supply temperature setpoints, qualitatively consistent with the context-dependent heating policy observed at Bülsweg (Figure 11).

A reward computation bug was active throughout our experimental run: denoiser coefficients were applied to raw feature values instead of z-score normalized inputs, and the baseline normalization

step was omitted. Internal rewards ranged from −81 to −15 kWh instead of the correct dimensionless 0.48–1.55 range. Despite this, a counterfactual analysis demonstrates that the impact on action selection was marginal in the operating temperature range: the missing baseline normalization removes a factor dependent only on outdoor temperature T, not on action Yb. At a fixed context T*, dividing by a positive scalar does not change argmin over Yb. At T = 8–10°C, the live system's actions agreed with the correct-reward optimum to within 0.08–0.25°C. Figure 26 compares the buggy GP surrogate against the reconstructed correct GP across three representative outdoor temperatures. The two reward landscapes exhibit highly similar shapes, and their minima — i.e. the heating actions chosen by Bayesian Optimization — coincide closely. This visual evidence confirms quantitatively that the bug shifted the reward scale but preserved its functional structure with respect to the action variable, leaving the optimization trajectory effectively intact.

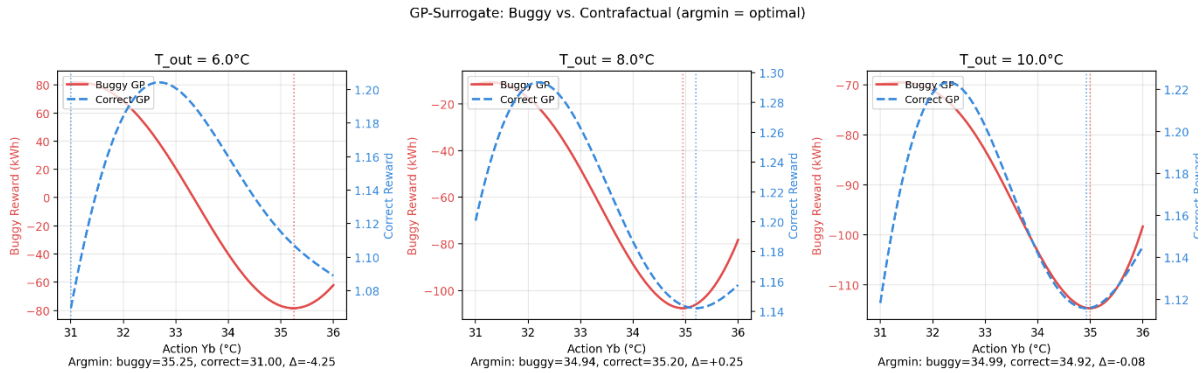

Figure 26: Counterfactual GP analysis: buggy GP (red, left axis) vs. correct GP (blue dashed, right axis) at T_out = 6, 8, 10°C. At T = 8 and 10°C, argmin values agree to within 0.08–0.25°C. At T = 6°C, deviation of 4.25°C is a domain boundary artefact. Bug had marginal practical impact in the core operating range.

Comfort: static comfort score 0.24 vs. adaptive comfort score of 0.28

Using equation (1.1). The calculation was based on indoor temperature measurements from each office at Sprint. For each qualifying sensor reading, we computed the deviation from the comfort thresholds: 23 °C during the day and 20 °C at night. These deviations were averaged per office and then aggregated across all offices to estimate the building-level comfort.

The resulting comfort scores were binned for visualization (Figure 17, right), and Gaussian distributions were fitted to both histograms. While the mean comfort score increased marginally from 0.24 (static) to 0.28 (adaptive), the variability of the distribution decreased substantially: the standard deviation dropped from 0.37 to 0.05, and the maximum observed comfort violation decreased from 2.0 to 0.38. The static reference period exhibited a long tail of severe deviations, whereas the adaptive phase produced a tightly concentrated distribution. This suggests that the adaptive controller does not merely preserve average comfort, but actively suppresses comfort extremes.

We note that the comfort term in the BO reward function (with comfort_weight = 1.0) was not affected by the denoiser bug discussed in Section 3.2.3.5.

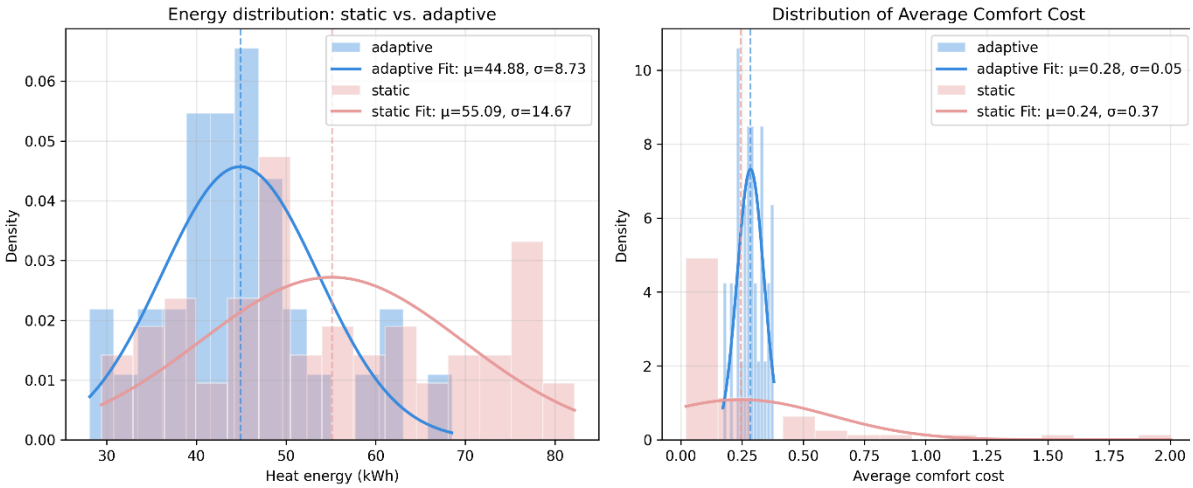

Figure 27: Distribution of daily heat energy (left) and comfort score (right) for static (grey) and adaptive AHA (blue) phases. Energy distribution shifts leftward and narrows. Comfort score: static shows right tail to 2; adaptive tightly concentrated around 0.28.

3.2.3.7. Discussion Energy Net Savings

The evaluation methods yield consistent results, with saving estimates progressively decreasing as seasonal confounders are controlled. Table below summarises the four estimates obtained under increasingly strict experimental control:

Method	Confounders controlled	Saving
Raw mean difference	None	17.6%
HDD normalization (bootstrap, T_base = 22°C)	Outdoor temperature, range restricted	22.4% ± 0.7%
Control period comparison (vs. Feb/March 2024 and 2025)	Season, temperature, solar irradiance	7.8–10.3%
Multivariate cross-validation (E~T+ Solar + Wind, R^2 = 0.83)	Weather, with seasonal bias correction	5.1–9.5%

Each method refines the previous by progressively isolating the AHA-attributable component from background variability. The convergent estimate from the two methods that include seasonal control (rows 3 and 4) places the **net AHA effect at NEST Sprint in the range of 5–10%.**

Thermal comfort. Despite a marginally higher mean comfort score in the adaptive phase (0.28 vs. 0.24), the algorithm eliminated comfort extremes: the standard deviation decreased by a factor of seven (0.37 → 0.05) and the maximum violation by a factor of five (2.0 → 0.38). This indicates that AHA modulates heating more stably than the static reference, simultaneously avoiding both wasteful over-heating and acute comfort deficits.

Comparison with Bülsweg. The Sprint net saving of 5–10% is broadly consistent with, and at the upper end somewhat exceeds, the Bülsweg HDD-normalized results (4.1% in Winter 2024/25). Several structural differences between the two deployments may contribute to this:

- The Sprint deployment had no Yc hardware freeze, so the full 1D action space along the Yb dimension was available throughout the experiment.
- The ceiling heating system at Sprint may respond more directly to Yb adjustments than the underfloor system at Bülsweg, given different thermal time constants.

A direct numerical comparison should therefore be treated with caution, even where the Sprint result appears nominally larger.

3.3 Field Test / Benchmark

A central hypothesis of the project proposal was that the proposed adaptive control approach should be applicable across diverse building types and heating systems without architectural modification. The two field deployments — Bülsweg and NEST Sprint — were deliberately selected to span a wide structural and operational range, providing an empirical test of this hypothesis.

The two installations differ along nearly every relevant dimension: building type (multi-family residential vs. office), heat source (district heating vs. heat pump), heat distribution (underfloor vs. ceiling heating), thermal mass (high vs. medium), occupancy pattern (continuous residential vs. business-hour office), and evaluation season (intra-seasonal vs. inter-seasonal). Despite these differences, the same algorithm — applied with identical hyperparameter configuration (length-scales, comfort weight, number of safe points) — produced consistent qualitative and quantitative behavior at both sites. Energy savings were of comparable magnitude (Bülsweg: 4–6%, Sprint: 5–10% after seasonal correction), the Bayesian Optimization showed stable adaptive convergence in both cases, comfort was preserved or improved, and the learned policies exhibited statistically significant context-dependence in the physically expected direction. No site-specific tuning, no building model, and no domain-specific feature engineering were required.

The minimal data requirements — three measured quantities (room temperature, outdoor temperature, heating energy) that are available in any modern building automation system — further support the transferability claim. The next phase will extend this evaluation to single-family houses with heat pumps as a third building typology, deliberately targeting the residential segment that is most relevant for broad market deployment.

Aspect	**Bülsweg**	**Sprint**
Building type	Multi-family residential	Multi office research
Heat source	District heating	Heat pump
Heat distribution	Underfloor	Ceiling
Thermal mass	High	Low
Occupancy	Continuous	Business hours
BO iterations	60	30
Hyperparameter set	Default	Default
Energy saving	4 – 6 %	5 – 10 %
Comfort outcome	Maintained	Maintained

4 Summary and Conclusions

4.1 Energy Performance and Savings Potential

The adaptive heating curve, controlled by Contextual Bayesian Optimization, was field-tested at two structurally distinct buildings. Both deployments yielded measurable energy savings, although each was subject to specific deployment constraints that likely reduced the achievable potential.

Bülsweg (multi-family residential, district heating, underfloor heating) Four complementary evaluation methods were applied:

- Linear Regression: 5.7% savings
- HDD normalization: 4.1% savings
- Machine Learning virtual control: 1.6% savings
- Digital Twin DID: not statistically significant (95% CI: –39.7 to +17.5 and –30.9 to +24.1)

The HDD-normalized estimate of 4.1% is the most reliable, as it corrects for weather variability without strong model assumptions. The Digital Twin showed discrepancies likely due to limited extrapolation beyond its training range. Importantly, this result is likely conservative: the adaptive configuration enforced a minimum inlet temperature of 30°C, whereas the previous static curve allowed temperatures to drop as low as 20°C. Analyses of high-temperature inefficiencies (Figure 17) further suggest up to 16% additional energy use under mild conditions, indicating that the present numbers reflect the inefficiency of an over-constrained optimization range rather than the algorithm's actual potential. Future experiments should extend the inlet temperature range to 20–45°C.

NEST Sprint (office unit, heat pump, ceiling heating) The static reference period (October–November 2025) and the adaptive phase (February–March 2026) fall in different seasons, requiring a dedicated seasonal correction. Two historical control periods (February–March 2024 and 2025, both verified as pure static operation) quantified the seasonal baseline shift at 10–12%. After removing this confounder, the net AHA-attributable saving converges to 5–10% across two methodologically independent approaches (direct comparison at common reference temperature; multivariate cross-validation), with a central value of approximately 7.5%. A reward computation bug — missing baseline normalization in the denoiser — was active throughout the deployment but, as discussed in Section 4.3, did not materially affect action selection in the operating temperature range.

4.2 Thermal Comfort

Thermal comfort was evaluated through a daily comfort score based on indoor temperature deviations from 23°C (day) and 20°C (night).

At **Bülsweg**, comfort levels remained nearly unchanged (static: 0.24, adaptive: 0.23), confirming that the adaptive strategy preserves occupant comfort in residential settings, with comparable comfort extremes, the standard deviation increased by 42 %: (0.14 → 0.20).

At **NEST Sprint**, while the mean comfort score increased marginally (0.24 → 0.28), the adaptive controller strongly suppressed comfort extremes: the standard deviation decreased by 86% (0.37 → 0.05) and the maximum violation by 81% (2.0 → 0.38). The static reference period exhibited a long tail of severe deviations, whereas the adaptive phase produced a tightly concentrated distribution. This suggests that the algorithm not only preserves average comfort but actively suppresses occupant-relevant comfort extremes. This represents a clearer comfort improvement than at Bülsweg, where mean and variability remained largely unchanged.

Across both deployments, thermal comfort was maintained.

4.3 Convergence and Robustness of the Learning Model

The Bayesian Optimization model demonstrated stable convergence after approximately 30 – 40 datapoints, beyond which additional data yielded no substantial improvement in predicted outcomes. This indicates a robust learning process within the explored context–action space.

However, robustness tests revealed an asymmetry in sensitivity:

- The model tolerated strong positive outliers (factor 3 distortion) without losing predictive accuracy.
- Conversely, negative outliers (factor 0.1 distortion) severely degraded performance, flattening predictions and leading to incorrect recommendations (e.g., excessive inlet temperatures at mild outdoor conditions).

This behavior suggests that underestimating rewards (negative bias) is more harmful than overestimation, emphasizing the need for regularization and noise handling strategies in future iterations of the BO framework.

4.4 Overall Conclusion

Overall, the adaptive heating curve guided by Bayesian Optimization demonstrates promising energy savings, validated across two building types and heating systems. At Bülsweg (multi-family residential, underfloor heating), HDD-normalized analysis confirms savings of approximately 4–6% with maintained thermal comfort. At NEST Sprint (office unit, ceiling heating), control-period comparison isolates a net AHA-attributable saving of 5-10% with an 86% reduction in comfort variability. Deployments across these two distinct typologies confirm that the approach generalizes across building types, with consistent energy savings and maintained or improved thermal comfort in both cases. These results support the transferability of the AHA framework as a scalable and data-efficient alternative to static heating curve control.

5 Outlook

Our industrial partner Co4 / Lippuner AG has expressed strong interest in continuing the collaboration:

"We have great interest in continuing this work. Our goal is to provide customers with a tool that not only monitors heating systems but actively optimizes their operation. Savings of around 5% are already significant. For future iterations, integrating predictive weather forecasting would be an exciting addition."

This perspective underlines the practical relevance and commercial potential of the adaptive heating control system. The next development phase, partially supported by internal resources beyond the scope of the present BFE project, focuses on the priorities outlined below.

Automation and observability. A fully automated deployment pipeline became operational during Winter 2025/26. Continuous deployments can be created via cron job, drawing a daily sample, refitting the Gaussian Process, and automatically setting the GP minimum as the new heating curve parameter. Basic logging verifies setpoint write-back consistency. This had been a recurring issue during Winter 2024/25, where setpoints were occasionally not propagated to the controller. Operating experience revealed, however, that fully automated operation cannot run unattended: complementary monitoring is needed, particularly during the early phase of each deployment, to confirm that the fitted reward landscape behaves as expected. With such early-stage tracking in place, both the denoiser misapplication at Sprint and an oversized GP length-scale identified in an earlier run would have been detected within the first few iterations rather than only in post-hoc analysis. Future work will therefore focus less on

adding automation and more on adding observability: real-time diagnostics of reward signals, GP fits, and convergence trajectories.

Better initialization. The current heating curve initialization relies on rough rule-of-thumb estimates for the four anchor points. Better-grounded prior estimates — derived from building characteristics, historical operation data, or system identification methods — could meaningfully accelerate convergence and reduce the number of samples needed before the BO reaches a useful operating point. A systematic study of how initialization quality affects deployment performance is planned, with the aim of providing a transferable initialization procedure for new sites.

Coupled parameter shifts. The current implementation optimizes a single heating curve parameter (Yb) while keeping the others fixed, in order to maintain a tractable 1D action space. Independently shifting multiple parameters quickly becomes sample-inefficient, as the search space grows multiplicatively. A promising extension is the coupled shifting of multiple curve parameters under a fixed proportionality — e.g., shifting Yb such that the ratio Yb/Yc remains constant. This approach, inspired by recent work on outdoor-temperature-dependent heating curve shifts [7], allows the algorithm to explore a richer family of curves while keeping the effective dimensionality of the optimization low. Combined with smarter initialization, this could unlock additional savings without sacrificing the data efficiency that is the core advantage of the current approach.

Maximum potential analysis. The current implementation prioritizes interpretability and data efficiency over exhaustive parameter exploration. Using simulation studies, we will explore the full parameter space of the heating curve to estimate its theoretical optimization potential and benchmark our current lightweight implementation against it. This study will also inform the practical limits of coupled parameter shifts (above) and provide quantitative targets for the full-season deployments.

Methodological improvements. Several extensions are planned: integrating wind and solar gains directly into the Gaussian Process model as additional context variables; evaluating whether the fixed 24-hour update cycle remains optimal or should be adapted to system dynamics; refining the exploration–exploitation balance and acquisition function step sizes to accelerate convergence; and systematically optimizing kernel functions and the relative weighting of energy versus comfort scores (currently fixed at 1.0 each), allowing the model to adapt to building-specific characteristics.

Robustness. Building on the asymmetric sensitivity observed at Bülsweg (Section 4.3), filtering and robust optimization techniques such as reward clipping and outlier detection will be implemented to ensure stable performance under imperfect data conditions.

6 National and international cooperation

Lippuner AG brings decades of expertise in energy and building technologies, with a strong commitment to advancing these fields through innovation. They provide the necessary infrastructure and practical experience to integrate and monitor our adaptive Heating Tuner within their systems. Additionally, they facilitate testing our algorithm on the Bülsweg building, which they operate, while sharing essential data and proactively addressing sensor failures.

7 Data management plan and open access/data/model strategy

In alignment with our commitment to open access and the requirements for project co-funding by the Swiss Federal Office of Energy (SFOE), we have developed a comprehensive data management plan. Our approach includes making parts of the algorithm's code, datasets, and digital twin models publicly accessible through the Renku platform [6].

By early next year, we plan to release the developed Python interface for direct interaction with EnergyPlus models. This interface, designed for seamless integration and ease of use, will be made available independently of ongoing collaboration with Lippuner AG, enabling broad usability and accessibility. The current development version can be accessed here: https://renkulab.io/p/michael.locher/adaptive-heizung-anpassung [6].

Furthermore, insights, results, and developed programs from the project will be published on the ARAMIS platform to ensure transparency and knowledge sharing within the community. This effort underlines our dedication to fostering open research and reproducibility while supporting innovation in energy efficiency and digital twin technologies.

8 References

ID	Reference description
1	Krause, Andreas and Ong, Cheng Soon (2011). Contextual Gaussian process bandit optimization, https://dl.acm.org/doi/10.5555/2986459.2986732.
2	Sui, Yanan and Gotovos, Alkis and Burdick, Joel and Krause, Andreas (2015). Safe Exploration for Optimization with Gaussian Processes, https://proceedings.mlr.press/v37/sui15.html.
3	Marcello, Fiducioso and Sebastian, Curi and Benedikt, Schumacher and Markus, Gwerder and Andreas, Krause (2019). Safe Contextual Bayesian Optimization for Sustainable Room Temperature PID Control Tuning, https://arxiv.org/abs/1906.12086
4	Sandmeier, E., Lobsiger-Kägi, E., Marek, R., Tomić, U., & Kälin, S. (2020). *2000-Watt-Gesellschaft leben: Reduktion des End-Energieverbrauchs durch Verhaltensänderungen – Nutzerinterventionen im Hüttengraben-Areal*. Schlussbericht vom 23. Dezember 2020. Eidgenössisches Departement für Umwelt, Verkehr, Energie und Kommunikation UVEK, Bundesamt für Energie BFE, Sektion Energieforschung und Cleantech. Available: https://www.aramis.admin.ch/Default?DocumentID=67511&Load=true, Chapter 4 Ergebnisse und Diskussion
5	PrimoŽ PotoČnik & Edvard Govekar (2019) Adaptive optimization of heating curves in buildings heated by a weather-compensated heat pump, Science and Technology for the Built Environment, 25:10, 1380-1393, DOI: 10.1080/23744731.2019.1616984. Available: https://doi.org/10.1080/23744731.2019.1616984
6	Data management Platform from SDSC: https://renkulab.io/
7	Potthoff, U., Brudermueller, T., Hopf, K., & Wortmann, F. (2025). Optimization of heating curves for heat pumps in operation: Outdoor temperature ranges for energy-efficient heating curve shifts. *Applied Energy*, *389*, 125455. https://doi.org/10.1016/j.apenergy.2025.125455
8	NEST Sprint, accessed April 2026

Appendix

- Appendix A: Bayesian Optimization and Gaussian Process Regression for Control Tuning
- Appendix B: Heating Degree Day
- Appendix C: Simulation
- Appendix D: Optimal Parameters for score function
- Appendix E: Data management plan and open access/data/model strategy
- Appendix F: Overview Bülsweg
- Appendix G: Data driven digital twin
- Appendix H: Pseudocode of the bootstrapped Difference-in-Differences algorithm

Appendix A: Bayesian Optimization and Gaussian Process Regression for Control Tuning

We defined the factors of our score function in Chapter 1.2. This function allows us to determine the scores associated with the optimization problem. However, we lack direct knowledge of how this function generally behaves or how it specifically applies to a given building. To understand and describe this function in order to optimize it later, we need to perform measurements and calculate the associated scores. Each measurement corresponds to a single data point. Once we have collected enough measurements, we can use these data points to create a surrogate model, see chapter 3.1.3.

We model our score function using GPR, which is particularly suited for this purpose due to its ability to approximate the true score function of a building with minimal prior knowledge. In the case of Contextual Bayesian Optimization, using GPR results in a three-dimensional surface, as illustrated in figure 27 and detailed in [1].

The goal of the next step is to use BO to identify the optimal action – in our case, the inlet temperature – that minimizes scores. BO proceeds as follows: since the context z is an external, uncontrollable variable, we fix the context at each step. In our example, we use the previous day's context value, assuming that the temperatures on two consecutive days do not differ significantly. By doing this, the optimization problem is reduced to a one-dimensional curve. This one-dimensional curve now describes the relationship between scores and the inlet temperature. This curve, along with its confidence intervals, is determined by the surrogate model determined by GPR.

Based on the newly acquired measurement, BO identifies the next best candidate for the inlet water temperature. This temperature is then set as the new value for the heating curve, and the next measurement is performed. Initially, each measurement will result in significant changes to the three-dimensional surface and, consequently, to the inlet water temperature. However, as more measurements are taken, these changes will decrease, eventually leading to convergence. This means we can approximate the effective underlying score function of an individual building with high accuracy.

In summary:

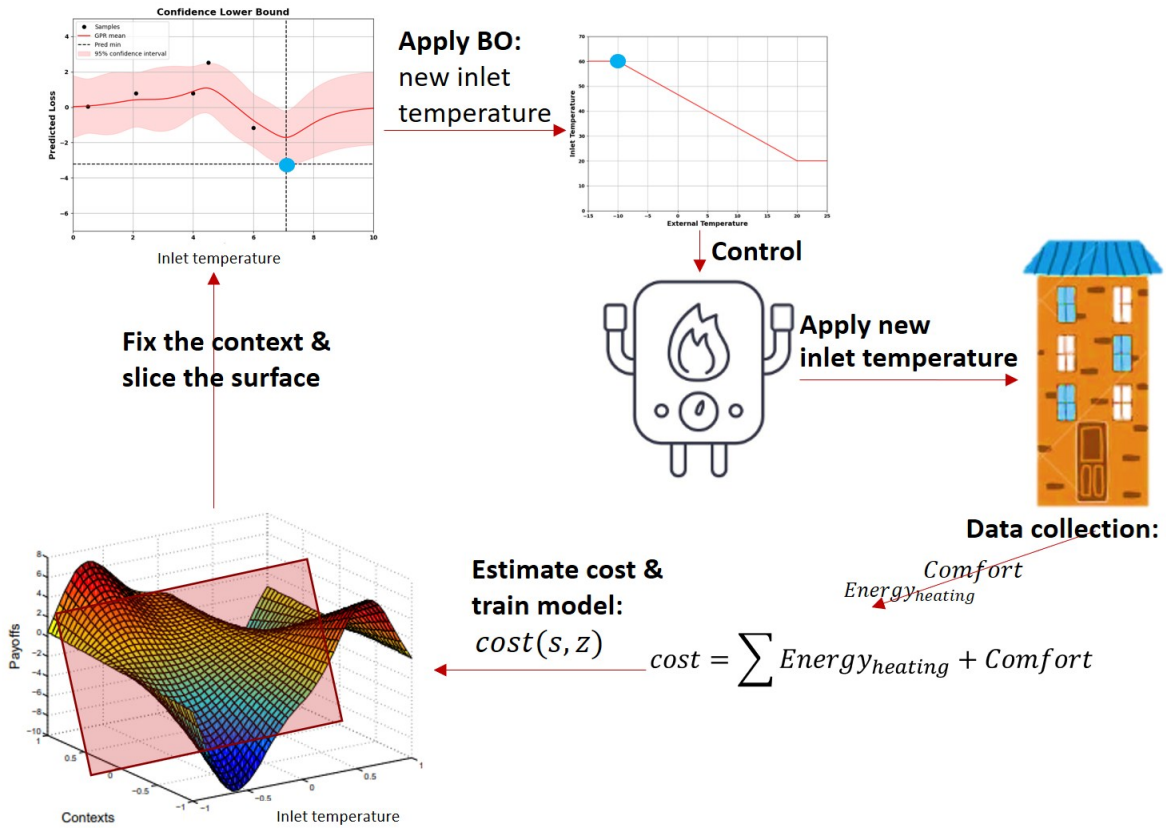

Figure 27: Depicted is the workflow for the contextual optimization of AHA. The experiment is initiated by setting the inlet temperature of the heating system, which leads to the Data Collection phase to gather observations needed to determine the score function. The scores are estimated, and the GPR is trained on the gathered information to obtain the 3D surface structure representing the score of the examined building. This provides the estimated scores (mean and variance) for each context and inlet temperature. Next, we fix the context based on past values and slice the surface, resulting in a 1-dimensional curve that depends only on the inlet temperature. We then apply BO to propose the next inlet temperature that minimizes the scores under the fixed context.

Appendix B: Heating Degree Day

The HDD method is a straightforward approach that normalizes heating energy based on heating degree days, providing a single value – the normalized heating energy consumption – for comparison. This allows for direct evaluation of energy use before and after an intervention.

Heating degree days reflect the climate-related heating energy demand by calculating, for each heating day (when the daily mean temperature is ≤ 12°C), how much the measured outdoor air temperature deviates from the desired indoor temperature of 20°C. The HDD 12/20 metric represents the sum of temperature differences from 20°C across all heating days within a specific period.

The HDD method enables a simple adjustment for outdoor air temperature by dividing the total heating energy consumption by the number of heating degree days. The resulting value indicates the average heating energy consumption per degree day, which can be compared across different periods (e.g., before and after an intervention). Using this approach, HDD values are determined for each cluster.

The main advantage of this method lies in its simplicity, as it distils heating energy performance into a single, easily comparable metric.

$$HDD = \sum_{k=0}^{n} \max(0, \mathrm{T_{set}} - \mathrm{T_{outside}})$$

$$normalized\ Heatingenergy\ = \frac{\sum_{k=0}^{n} Heatingenergy_k}{HDD}$$

Appendix C: Simulation

We have a complete setup to simulate EnergyPlus models using various heating curve models. For this, we have EnergyPlus models for three buildings: Bülsweg, NEST NEST, and NEST Sprint. We have also developed an interface and a data processing pipeline for our machine learning models. For these buildings, we can generate an unlimited number of weather data sets and create numerous simulation scenarios. Additionally, real data is available for these buildings. For Bülsweg, the heat output can be estimated through digitally readable heat meters. Each residential unit is also equipped with a multifunction sensor to measure VOC concentration, room temperature, and relative humidity throughout the project duration. Furthermore, we have access to weather data via a weather station in Vaduz. The outdoor sensor of the heating system is used to determine the ambient temperature. An overview of the AHA program can be found in figure 28. The table below explains the key classes and the source code: adaptive heizung anpassung · Project in @michael.locher · Renku

Class	Description
SimulationWrapper	Interface to run Energyplus models containing specification of building / heating coefficients and variables used for our project which are extracted from the energyplus output.
SimulationDriver	Reading and processing the EnergyPlus output data so that it can be used for our machine learning models. Additionally, it includes the Controller class, which integrates all components (models, algorithms, and data storage) to perform effective sampling.
Denoiser	Initial measurements based on static heating curves indicate that the variance in heating requirements for each outdoor temperature can be quite large. This is primarily due to the impact of solar radiation. The effect of solar radiation is estimated and subtracted from the heating energy. The remaining heating energy is then standardized using a reference curve.
BoModel	The BO Heating model keeps track of all the samples and combines the various elements composing the model (the algorithm, the surrogate model, the heating curve)
SurrogateModel	The surrogate model is used to predict the score of a given action in a given context. We use the Gaussian Process Regression.
HeatingCurve	The HeatingCurve class models the parameter space of the heating curve, it contains the domain and the initial safe set.
UCBAlgo	We use the upper confidence bound to estimate the most promising next action using Gaussian Process Regression as surrogate model.

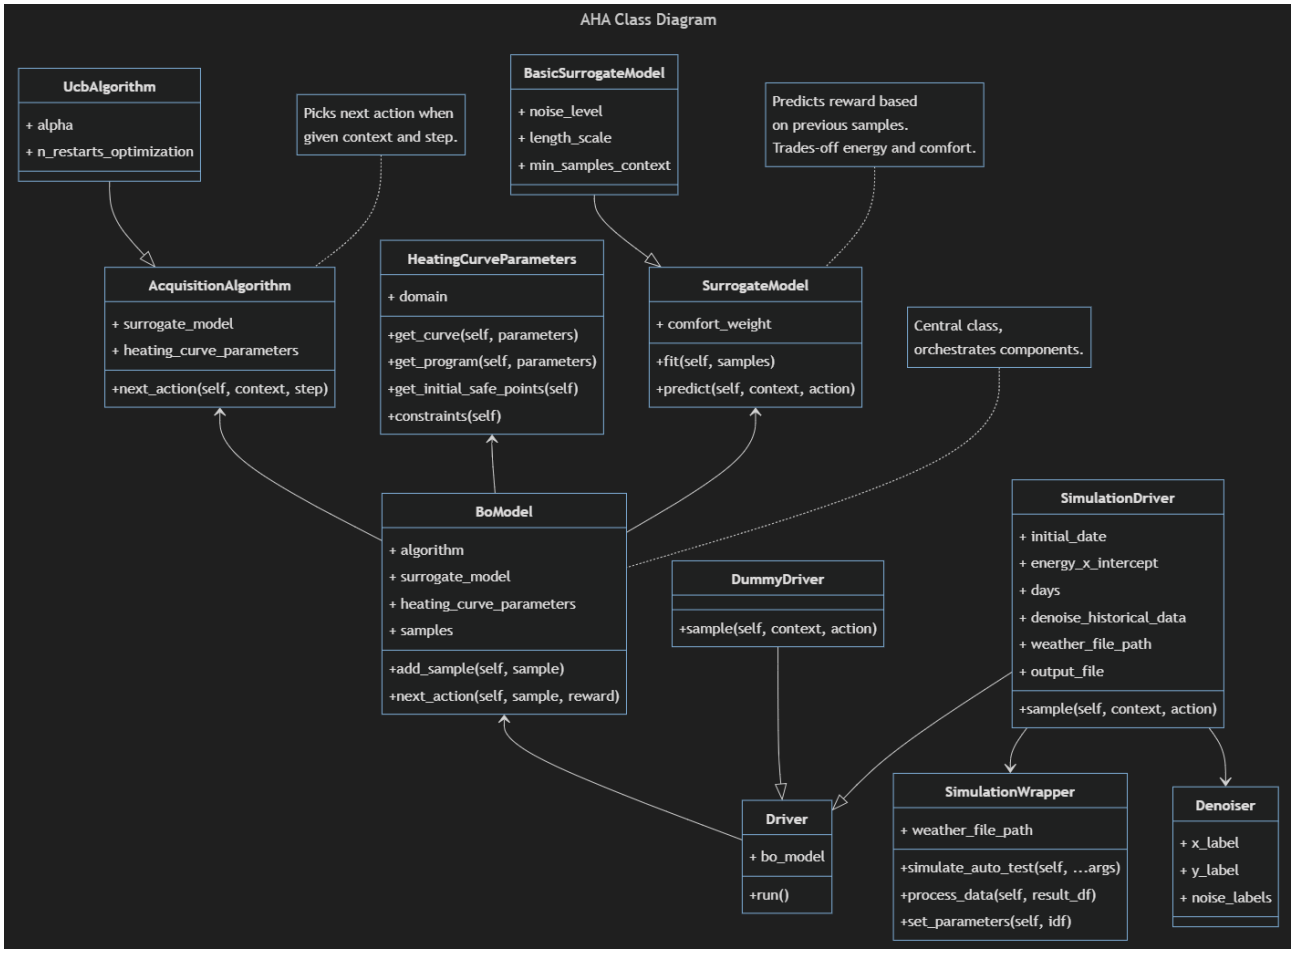

Figure 28: AHA Class Diagram displaying the different classes, attributes, methods and dependencies.

Appendix D: Optimal Parameters for score function

The definition of the score function in Chapter 1.2 allows us to experiment with the weighting between comfort and heating energy:

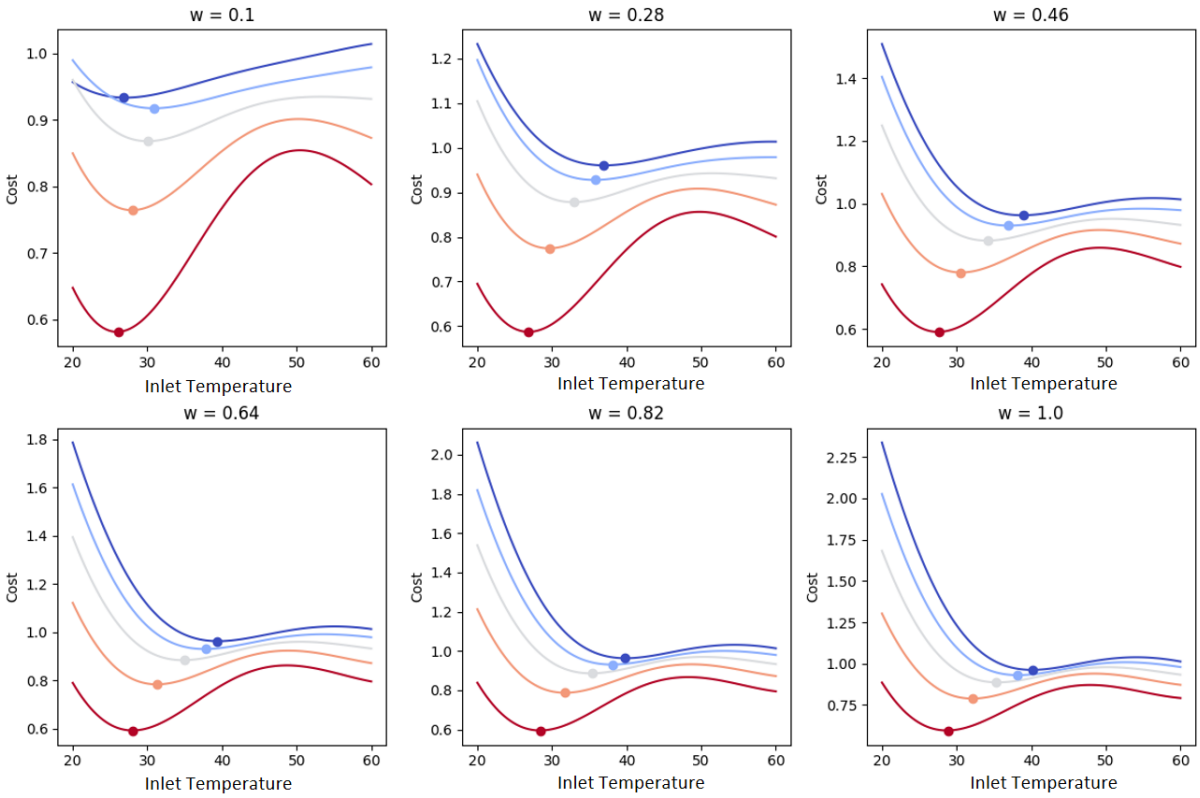

Figure 29: The simulations for NEST clearly demonstrate how adjusting the weighting between heating energy and comfort allows for additional energy savings. Of particular interest is the range between a weighting of 1, which emphasizes minimizing comfort score, and a slack at a weighting of approximately 0.8.

Appendix E: Data management plan and open access/data/model strategy

In alignment with our commitment to open access and the requirements for project co-funding by the Swiss Federal Office of Energy (SFOE), we have developed a comprehensive data management plan. Our approach includes making parts of the algorithm's code, datasets, and digital twin models publicly accessible through the Renku platform [6].

By early next year, we plan to release the developed Python interface for direct interaction with EnergyPlus models. This interface, designed for seamless integration and ease of use, will be made available independently of ongoing collaboration with Lippuner AG, enabling broad usability and accessibility. The current development version can be accessed here: adaptive heizung anpassung · Project in @michael.locher · Renku [6].

Furthermore, insights, results, and developed programs from the project will be published on the ARAMIS platform to ensure transparency and knowledge sharing within the community. This effort underlines our dedication to fostering open research and reproducibility while supporting innovation in energy efficiency and digital twin technologies.

Appendix F: Overview Bülsweg

The Bülsweg building consists of three 2.5-room apartments and six 3.5-room apartments. Bülsweg is connected to the district heating network of the Buchs waste incineration plant, and all units are heated through underfloor heating. Heat consumption per apartment can be regularly monitored via digital heat meters. During the project, each apartment will also be equipped with a multifunction sensor to measure

VOC concentration, indoor temperature, and relative humidity. Although a photovoltaic system is installed on the building's roof, we do not have access to its data. Instead, we rely on data from local weather stations – specifically, the Vaduz station – to determine solar irradiation. The outdoor sensor of the heating system provides information about the ambient temperature.

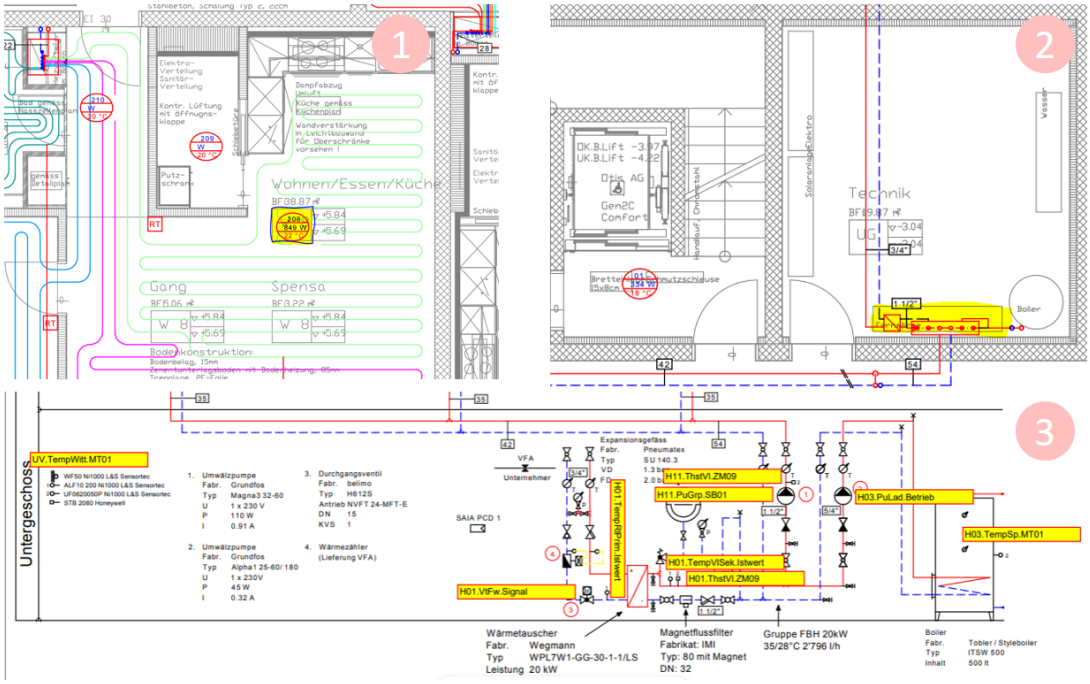

Figure 30: Subfigure 1 shows the floor plan with underfloor heating. Each room has an individual heating loop, which is configured to maintain the intended setpoint temperatures when the room thermostat is set to its default setting. For example, the reference room "Living/Dining/Kitchen" is heated to 22°C when the thermostat is set to level 3. Residents can fine-tune the flow through the room thermostat, allowing them to overheat or underheat the room according to personal preferences, thus adjusting the effective heating output. The heating power is supplied by the district heating system from Grabs and is transferred to a heat distributor in the basement, which then distributes the heat throughout the entire house. All key parameters for determining comfort and heating performance can be accessed for Bülsweg.

Key	Description	Dim	Freq	Installation
H01.TempVlSek.MT01	Secondary supply temperature	°C	15 min	Heating room
H01.TempRlPrim.MT01	Primary return temperature	°C	15 min	Heating room
H01.VtFw.VA01	Forward flow rate	m³/h	15 min	Heating room
H03.TempSp.MT01	Boiler temperature	°C	15 min	Heating room
UV.TempWitt.MT01	Outdoor temperature	°C	15 min	Weather sensor
H11.Anlage.Sollwert	System setpoint value	°C	15 min	Heating room
H11.Anlage.Errechnet	Calculated system output	°C	15 min	Heating room
H11.Anlage.Kurve_Yb_Empa	Heating curve setting Yb Empa	°C	15 min	Heating room
H03.PuLad.SB01	Loading pump command	On/Off	15 min	Heating room
H11.PuGrp.SB01	Pump group switching command	On/Off	15 min	Heating room
H11.PuGrp.Schaltung	Pump group switching status	OnOff	15 min	Heating room
H11.PuGrp.Stunden	Pump group operating hours	h	15 min	Heating room
UV.TempWitt.VA01	Outdoor temperature (analog value)	°C	15 min	Weather sensor

Key	Description	Dim	Freq	Installation
global_solar	Global solar radiation	W/m²	hourly	MeteoSwiss station Vaduz
diffuse_solar	Diffuse solar radiation	W/m²	hourly	MeteoSwiss station Vaduz
windspeed	Wind speed average	m/s	hourly	MeteoSwiss station Vaduz
precipitation	Precipitation sum	mm	hourly	MeteoSwiss station Vaduz
wind_direction_mean_deg	Mean wind direction (hourly)	°	hourly	MeteoSwiss station Vaduz
wind_direction_10min_deg	Wind direction (10-minute value)	°	10 min	MeteoSwiss station Vaduz
relative_humidity	Relative humidity	%	hourly	MeteoSwiss station Vaduz
sunshine_duration	Sunshine duration	min	hourly	MeteoSwiss station Vaduz
mean_temperature	Air temperature at 2m	°C	hourly	MeteoSwiss station Vaduz
dew_point	Dew point temperature at 2m	°C	hourly	MeteoSwiss station Vaduz
date	Timestamp	-	hourly	MeteoSwiss station Vaduz

Apartment	Key	Description	Dim	Freq	Installation
WHG1 EG Ost	temperature	Indoor temperature	°C	irregular	Entry living space, 1.5 m
	humidity	Indoor humidity	%	irregular	Entry living space, 1.5 m
	co2	CO_2 concentration	ppm	irregular	Entry living space, 1.5 m
	heat_electricity	Heating electricity consumption	kWh	irregular	Heat counter
WHG2 EG Mitte	temperature	Indoor temperature	°C	irregular	Entry living space, 1.5 m
	humidity	Indoor humidity	%	irregular	Entry living space, 1.5 m
	co2	CO_2 concentration	ppm	irregular	Entry living space, 1.5 m
	heat_electricity	Heating electricity consumption	kWh	irregular	Heat counter
WHG3 EG West	temperature	Indoor temperature	°C	irregular	Entry living space, 1.5 m
	humidity	Indoor humidity	%	irregular	Entry living space, 1.5 m
	co2	CO_2 concentration	ppm	irregular	Entry living space, 1.5 m
	heat_electricity	Heating electricity consumption	kWh	irregular	Heat counter
WHG4 1OG Ost	temperature	Indoor temperature	°C	irregular	Entry living space, 1.5 m
	humidity	Indoor humidity	%	irregular	Entry living space, 1.5 m

	co2	CO_2 concentration	ppm	irregular	Entry living space, 1.5 m
	heat_electricity	Heating electricity consumption	kWh	irregular	Heat counter
WHG5 1OG Mitte	temperature	Indoor temperature	°C	irregular	Entry living space, 1.5 m
	humidity	Indoor humidity	%	irregular	Entry living space, 1.5 m
	co2	CO_2 concentration	ppm	irregular	Entry living space, 1.5 m
	heat_electricity	Heating electricity consumption	kWh	irregular	Heat counter
WHG6 1OG West	temperature	Indoor temperature	°C	irregular	Entry living space, 1.5 m
	humidity	Indoor humidity	%	irregular	Entry living space, 1.5 m
	co2	CO_2 concentration	ppm	irregular	Entry living space, 1.5 m
	heat_electricity	Heating electricity consumption	kWh	irregular	Heat counter
WHG7 2OG Ost	temperature	Indoor temperature	°C	irregular	Entry living space, 1.5 m
	humidity	Indoor humidity	%	irregular	Entry living space, 1.5 m
	co2	CO_2 concentration	ppm	irregular	Entry living space, 1.5 m
	heat_electricity	Heating electricity consumption	kWh	irregular	Heat counter
WHG8 2OG Mitte	temperature	Indoor temperature	°C	irregular	Entry living space, 1.5 m
	humidity	Indoor humidity	%	irregular	Entry living space, 1.5 m
	co2	CO_2 concentration	ppm	irregular	Entry living space, 1.5 m
	heat_electricity	Heating electricity consumption	kWh	irregular	Heat counter
WHG9 2OG West	temperature	Indoor temperature	°C	irregular	Entry living space, 1.5 m
	humidity	Indoor humidity	%	irregular	Entry living space, 1.5 m
	co2	CO_2 concentration	ppm	irregular	Entry living space, 1.5 m
	heat_electricity	Heating electricity consumption	kWh	irregular	Heat counter

Appendix G: Machine Learning-Based Digital Twin

G.1 Architecture

The digital twin developed in this work is based on a **cascaded architecture** of machine learning models. The objective of this architecture is to predict the thermal and energetic dynamics of the building and its heating system in a physically consistent, data-driven manner. The single models are interlinked such that the output of one stage forms an input to the subsequent stage. This chaining of predictions reflects the causal dependencies between different physical quantities (supply temperature, return temperature, thermal balance, heat consumption, and indoor comfort).

G.1.1. Cascaded Model Structure

The cascaded approach was selected to reflect causal dependencies in the heating system:

- The analytical inlet temperature forms the starting point, reflecting control setpoints.
- Predicted inlet temperature propagates to return temperature, ΔT, and subsequently to heating energy.
- Finally, heating energy and system state determine indoor temperature and thus comfort.

This **hierarchical design** ensures that predictions remain physically interpretable and robust against unrealistic extrapolation. For example, setting the inlet signal to zero leads to downstream predictions converging at significantly lower heat energy predictions

The digital twin consists of six interconnected models, see Figure 15:

Model	Target & Concept	Features	Performance (MAE / RMSE / R^2)
1. Analytical Inlet Temperature	Target: inlet_calculated Concept: Analytical heating curve calculation, baseline control signal	Heating curve parameters	–
2. Supply Temperature (TempVlSek)	Target: Measured supply temp TempVlSek Concept: Residual approach: • Base: Linear regression on inlet_calculated_old • Residuals: XGBoost for nonlinear effects	Base: inlet_calculated_old Residuals: time encodings, weather, hydraulics, control inputs	Static: MAE 0.851 / RMSE 1.438 / R^2 0.918 Dynamic: MAE 3.120 / RMSE 4.794 / R^2 0.364
3. Return Temperature (TempRlPrim)	Target: TempRlPrim Concept: XGBoost conditioned on predicted TempVlSek and hydraulic/control variables	Time encodings, hydraulics (VtFw, Pumpladung), TempVlSek, building temp	Static: MAE 0.684 / RMSE 1.053 / R^2 0.953 Dynamic: MAE 0.915 / RMSE 1.598 / R^2 0.924
4. Temperature Difference (ΔT)	Target: TempVlSek – TempRlPrim Concept: XGBoost proxy for heat transfer to building	Pump activity, TempVlSek, TempRlPrim, lagged ΔT, building temp	Static: MAE 0.311 / RMSE 0.523 / R^2 0.862 Dynamic: MAE 0.626 / RMSE 0.968 / R^2 0.599
5. Heating Energy (heat_electricity)	Target: Thermal energy consumption heat_electricity Concept: XGBoost approximating heat counter	Weather/climate drivers, hydraulics, time encodings, predicted states (TempVlSek, TempRlPrim, ΔT), autoregressive lags	Static: MAE 1.516 / RMSE 2.033 / R^2 0.876 Dynamic: MAE 2.158 / RMSE 3.048 / R^2 0.703
6. Indoor Building Temperature	Target: building_temp (avg indoor temp) Concept: XGBoost linking heating system + weather to indoor comfort	Predicted supply/return temps, energy use, weather drivers	Static: MAE 0.149 / RMSE 0.211 / R^2 0.868 Dynamic: MAE 0.135 / RMSE 0.190 / R^2 0.821

G.1.2. Feature Selection Process

The features for each model were selected using a systematic **feature group search procedure**. Features were first grouped into physically meaningful categories:

- Weather drivers: global_solar, diffuse_solar, sunshine_duration, windspeed, wind_direction_mean_deg, relative_humidity, TempWitt_analog, precipitation, dew_point
- Building state and thermal inertia: building_temp
- Heating curve setpoints and control actions: inlet_calculated_old, action
- System measurements (supply/return): TempVlSek, TempRlPrim, delta_T, heat_electricity
- Pumps and hydraulics: Pumpladung, VtFw, pump_active, pump_recent
- Calendar encodings (cyclical time variables): hour, dayofweek, hour_sin, hour_cos, dow_sin, dow_cos
- Autoregressive lags: building_temp_lag1, building_temp_lag2, building_temp_lag3

All combinations of feature groups were evaluated with an **autoregressive XGBoost regressor** on a train/test split of the static regime. The best-performing feature set for each target was selected to ensure that the models balance predictive performance with interpretability.

G.1.3. Sanity Checks

To ensure the physical consistency and robustness of the cascaded model, several sanity checks were conducted. These tests assess whether the model behaves plausibly under controlled perturbations of key input variables.

Sensitivity Analysis with Respect to Inlet Temperature

As a first check, we analyzed how the model responds to variations in the inlet temperature (inlet_calculated), which directly affects the overall energy balance. From a physical standpoint, when setting the inlet temperature to zero, we would expect the predicted supply temperature (TempVlSek) to also approach zero.

The results are as follows:

- **Baseline model:** predicts approximately **35 °C** for TempVlSek, reflecting the average inlet temperature observed during training.
- **Digital twin:** predicts around **7 °C**, which aligns more closely with the expected physical behavior under such extreme conditions and leads to a substantial reduction in predicted heating energy.

Perturbation Experiment

We conducted a controlled sensitivity test by adjusting the inlet temperature ±5 °C around the baseline, using a representative test window (01.02.2025 – 22.02.2025). Three scenarios were simulated with the cascaded model:

Scenario	Description	Total Predicted Energy (kWh)	Δ vs. Baseline (kWh)
Baseline	Original inlet temperature	2571.26	—
Inlet – 5 °C	Reduced inlet temperature	1918.15	–653.11
Inlet + 5 °C	Increased inlet temperature	3579.34	+1008.08

Based on the fundamental physics of heat transfer, one might expect a roughly symmetrical response in energy use when adjusting the inlet temperature, since heating demand is proportional to the temperature difference between the supply and the environment ($Q = m \cdot \Delta T$). In other words, decreasing or increasing the inlet temperature by the same amount should, in theory, result in comparable reductions or increases in energy consumption. However, the observed response is asymmetric: while lowering the

inlet temperature reduces energy demand, increasing it leads to a disproportionately larger rise in energy consumption.This asymmetry indicates that the model does not yet fully capture the linear physical relationship between inlet temperature and thermal energy demand.

Nevertheless, this sanity check confirms that the model reacts qualitatively correctly to changes in inlet conditions, higher inlet temperatures increase, and lower inlet temperatures decrease, the predicted heating energy, providing confidence in the general directionality of the model's response.

G.1.4. Performance Comparison

On the hold-out test set (dates: 2025-01-10 and 2025-01-23), the digital twin architecture outperforms a baseline, e.g. GAM model:

GAM baseline: RMSE = 3.22, MAE = 2.31, R^2 = 0.69

Digital twin: RMSE = 2.43, MAE = 1.81, R^2 = 0.82

Beyond these accuracy improvements, the digital twin also supports extrapolation to counterfactual scenarios, see Sanity Checks above.

This example highlights how the digital twin improves accuracy and yields more realistic extrapolations, enabling meaningful exploration of scenarios outside the training distribution.

Appendix H: Pseudocode of the bootstrapped Difference-in-Differences algorithm

This section provides the pseudocode of the bootstrapped Difference-in-Differences algorithm used to quantify the effect of the adaptive heating tuner compared to the static baseline while accounting for model bias and uncertainty. The procedure combines the outputs of two digital twins, DT1 (adaptive) and DT2 (static), to estimate the distribution of possible energy savings.

Step 1: Aggregate predictions

- Convert hourly predictions of both digital twins: DT1 = adaptive and DT2 = static, into daily totals: kWh per day
- Split into pre-period (before intervention) and after-period (after intervention).

Notation: $\overline{DT1}_{pre}, \overline{DT2}_{pre}, \overline{DT1}_{after}, \overline{DT2}_{after}$

Step 2: Estimate bias and variance

- For each DT and period, compute residuals against the real testbed TB:

$$r_{m,period} = TB_{period} - DT_{m,period} \text{ , } m \in (1,2)$$

- From these residuals, estimate:

$$\text{Mean bias } b_{m,period} = mean(r))$$

$$\text{Standard deviation } s_{m,period} = std(r)$$

- Interpretation: bias = systematic offset, std = model uncertainty
- The estimated bias and standard deviation are used to adjust the model noise, replacing the normal assumption $\varepsilon \sim N(0,1)$ with a more realistic noise model $N\left(b_{m,period}, s^2_{m,period}\right)$

Step 3: Bootstrap resampling procedure

1. For each bootstrap draw $b = 1, \ldots, B$: Resample days: sample daily values with replacement from each digital twin's daily series in pre and after

2. Inject model noise: For each resampled day, add Gaussian perturbation: $\varepsilon \sim N\left(b_{m,period}, s^2_{m,period}\right)$
3. Compute bootstrapped mean differences:

$$\Delta^{(b)}_{pre} = \overline{DT2}^{(b)}_{pre} - \overline{DT1}^{(b)}_{pre}$$

$$\Delta^{(b)}_{after} = \overline{DT2}^{(b)}_{after} - \overline{DT1}^{(b)}_{after}$$

4. Compute DID draw:

$$DID^{(b)} = \Delta^{(b)}_{after} - \Delta^{(b)}_{pre}$$

Step 4: Collect bootstrap distribution

- Repeat steps 1 – 3 for B iterations, in our case B = 5000
- Obtain the empirical distribution DID: $\{DID^{(1)}, \ldots, DID^{(B)}\}$

Step 5: Summarize effect

- Report the mean effect:

$$\overline{DID} = \frac{1}{B}\sum_{b=1}^{B} DID^{(b)}$$

- Report uncertainty as percentile-based confidence interval:

$$CI_{95\%} = [quantile_{0.025}, quantile_{0.975}]$$

Appendix I: Overview Sprint

NEST Sprint is a 167 m² office unit on the first floor of the NEST research platform at Empa Dübendorf, comprising 12 individual offices (rooms 171–186). The unit was opened on 16 August 2021 and built predominantly from reused materials and components. Note: rooms 174 and 175 have been combined into a single open office since 3 July 2023, and the partition wall between room 171 and the corridor was removed on 7 June 2024.

Heating energy is sourced from the NEST Medium Temperature Energy (MTE) grid (primary circuit 36/28°C) via a plate heat exchanger. Cooling energy is sourced from the NEST Low Temperature Energy (NTE) grid (primary circuit 10/20°C), likewise via a heat exchanger. Both heating and cooling energy are distributed to all rooms via a ceiling heating/cooling system (manufacturer: Barcolaire, type A11; heating capacity 5 kW total / 200 W per room; cooling capacity 2.5 kW / 400 W per room; mass flow 60 l/h per room). The system also provides thermal energy to a partial air conditioning unit (U10L1) for ventilation.

The supply temperature is governed by a weather-compensated flow temperature controller: the set-point is calculated as a function of the 24-hour mean outdoor temperature via an adjustable heating curve — the parameter optimized by AHA. Room temperatures are controlled individually via one proportional valve per room (U10Z7/U10Z8, Belimo CQ24A-SR). The building automation system (Beckhoff SPS CX5140) is accessible via the NEST API (OPC-UA, Modbus TCP/IP).

All key parameters for determining comfort and heating performance are accessible via the NEST API. Each office room is equipped with a ceiling-mounted multifunction sensor (Elsner Sewi KNX AQS/TH-D) measuring indoor temperature (°C), CO_2 concentration (ppm), relative humidity, and presence. Window contacts (Dragino LDS02 LoRaWAN) are installed per room. Total heating and cooling energy consumption is measured by dedicated heat meters at the zone level (U10Z1, Belimo 22PEM-1UC),

providing thermal power (kW), cumulative energy (kWh), volume flow (m³/h), and supply/return temperatures (°C). Total electrical consumption is monitored via a three-phase energy meter (U10E1). Outdoor temperature is available as a 1-hour running average from the on-site NEST weather station [8].

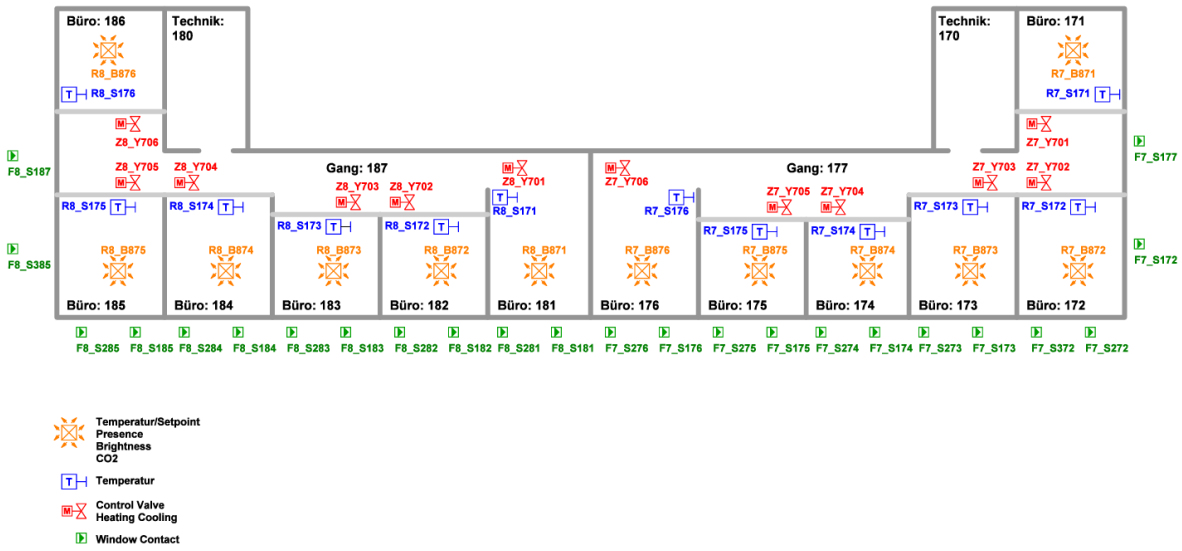

Figure 31: Floor plan NEST Sprint, 1st floor: room layout 171–186, ceiling heating/cooling elements, technical rooms 170 and 180 incl. installed sensors: window contacts, ceiling-mounted multifunction sensor.