

# Explaining AI Decisions: Towards Achieving Human-Centered Explainability in Smart Home Environments

Md Shajalal (✉)<sup>1,2</sup> , Alexander Boden<sup>1,3</sup> , Gunnar Stevens<sup>2,3</sup> , Delong Du<sup>2</sup> , and Dean-Robin Kern<sup>2</sup> 

<sup>1</sup> Fraunhofer Institute for Applied Information Technology FIT, Sankt Augustin, Germany

[md.shajalal@fit.fraunhofer.de](mailto:md.shajalal@fit.fraunhofer.de)

<sup>2</sup> University of Siegen, Siegen, Germany

<sup>3</sup> Bonn-Rhein-Sieg University of Applied Sciences, Sankt Augustin, Germany

**Abstract.** Smart home systems are gaining popularity as homeowners strive to enhance their living and working environments while minimizing energy consumption. However, the adoption of artificial intelligence (AI)-enabled decision-making models in smart home systems faces challenges due to the complexity and black-box nature of these systems, leading to concerns about explainability, trust, transparency, accountability, and fairness. The emerging field of explainable artificial intelligence (XAI) addresses these issues by providing explanations for the models' decisions and actions. While state-of-the-art XAI methods are beneficial for AI developers and practitioners, they may not be easily understood by general users, particularly household members. This paper advocates for human-centered XAI methods, emphasizing the importance of delivering readily comprehensible explanations to enhance user satisfaction and drive the adoption of smart home systems. We review state-of-the-art XAI methods and prior studies focusing on human-centered explanations for general users in the context of smart home applications. Through experiments on two smart home application scenarios, we demonstrate that explanations generated by prominent XAI techniques might not be effective in helping users understand and make decisions. We thus argue for the necessity of a human-centric approach in representing explanations in smart home systems and highlight relevant human-computer interaction (HCI) methodologies, including user studies, prototyping, technology probes analysis, and heuristic evaluation, that can be employed to generate and present human-centered explanations to users.

**Keywords:** Explainable AI (XAI) · Human-Centered XAI · Demand Forecasting · Machine Learning · Smart Home

---

This is the pre-print version of our accepted paper at the 2nd World Conference on eXplainable Artificial Intelligence (xAI2024), which will be held in Valletta, Malta in 17-19 July, 2024

## 1 Introduction

Due to advancements in sensor technology and machine learning (ML) over the past decades, smart home applications can now provide residents with the ability to monitor and control connected appliances via sensors [1]. These applications can even make decisions automatically using ML-driven techniques rather than relying on simple timetable logic. In the smart home energy domain, one notable energy-aware smart home application might be appliance-level energy-demand forecasting to make users more aware and help them optimize their energy consumption practices [1, 2]. Adjusting the heating system to provide a comfortable and healthy household and work environment based on predicting individuals’ thermal comfort preferences can be another fascinating energy-related smart home application [3–5]. Other applications also often utilize complex ML models to make decisions, such as human activity recognition within the home, identification of energy-intensive activities for different household tasks [6], fall detection and health monitoring [7], and energy optimization [1, 8].

AI-based applications in smart home systems are becoming increasingly popular as homeowners aim to enhance their living environment while reducing energy usage. Previous studies [1, 2, 8–10] have modeled energy demand forecasting in smart homes using AI techniques, including Deep Neural Networks (DNNs), Convolutional Neural Networks (CNNs), Auto-Encoders (AE), and Long-Short-Term Memory (LSTM). Classical Machine Learning (ML)-based predictive models have also garnered attention for predicting personal thermal comfort in indoor environments, thereby enhancing resident comfort [4, 11–17]. The complexity and opacity of these ML models often hinder their adoption in real-world scenarios due to difficulties in aiding users’ decision-making. Since ML models can be very complex, involving thousands to millions of model parameters (i.e., deep learning models), they are often referred to as *black-boxes*. Decisions from black-box models can be unintelligible and may surprise users with unexpected predictions. In such cases, users require explanations to comprehend the predictions. Recently, there has been significant interest in elucidating the decisions of ML models across various fields, including language processing [18], financial analytics [19], e-commerce [19], medicine and health [20], bioinformatics [21], and smart home applications [1, 8, 21]. This effort to clarify ML models’ decisions is referred to as eXplainable Artificial Intelligence (XAI).

XAI aims to develop AI systems that can provide clear explanations about the decision-making processes and the predicted decisions [22]. The “black-box” issue can lead to users’ mistrust and confusion about these technologies. To improve users’ trust and understanding, *Human-centered* explanations<sup>4</sup> can be a game-changer by providing clear and effective explanations to users, enabling them to troubleshoot issues and customize devices to suit their needs [23]. However, many XAI methods have been introduced and developed to explain the models’ decision-making procedures and the reasons behind specific predictions. Most of them are proposed to debug and improve models’ performance [1, 24, 25].

<sup>4</sup> Throughout the paper, “Human-centered XAI” and “user-centered XAI” are used interchangeably

In the context of smart home application scenarios, users in households are generally laypeople and may not have sufficient knowledge to understand technical explanations (i.e., even AI developers might struggle to understand explanations). Therefore, this paper advocates the need for easily understandable explanations for general users to make sense of predictions, which we refer to as “Human-centered” explanations.

While some studies have attempted to make models interpretable in the context of smart homes, most are focused on explaining predictions to improve performance and debug models [2, 26, 27]. As noted earlier, smart home users often lack the technical expertise necessary to understand many of the suggested explanations [25, 28, 29]. Additionally, research has shown that end-users have diverse perspectives when trying to make sense of smart home systems [30], and these perspectives can evolve over time. Various studies have demonstrated that current XAI techniques fail to produce human-centered explanations that assist general users in understanding the decision-making process and the reasons behind specific predictions [31, 32]. Given these findings, generating human-centered explanations for complex smart home applications is more challenging than might be initially assumed, especially considering the diverse backgrounds of general users.

Smart home systems encompass various energy consumption-related sub-tasks, including energy demand forecasting [1, 8], appliance-level consumption predictions [9], energy intensity identification for different household activities [6], and thermal comfort prediction [4, 5] for efficient heating systems. These systems are more complex than classical classification or regression tasks, and their collective outcomes contribute to the overall functionality of smart homes. However, the complexity of these systems can lead to decisions that are difficult for general users to understand, potentially hindering their adoption in real-world settings. This paper emphasizes the importance of human-centered explanations in smart home applications by reviewing the progress of state-of-the-art technical and human-centered XAI studies. We focus on two energy consumption-related application scenarios within smart home settings: energy demand forecasting and the prediction of personal thermal comfort preferences for smart heating systems. We conduct experiments by applying deep neural network-based energy forecasting and ML models on benchmark datasets to model thermal comfort preferences. To uncover the complexity of the predictions, we apply current prominent XAI methods to identify facts by presenting the explanations in various forms.

We present explanations generated by multiple XAI methods and analyze their understandability. We then identify the associated challenges that must be considered when generating human-centered explanations. The need for user-friendly explanations in these contexts illustrates the challenges in understanding complex decisions. Finally, we highlight several HCI methodologies that could be beneficial in achieving human-centered XAI in smart home applications. The contributions of this paper are threefold:

- We present state-of-the-art XAI methods, discuss progress towards human-centered XAI and highlight research gaps that hinder their immediate application in smart home systems.
- Through careful analysis of experimental results based on prominent explainability methods in two sub-tasks, we argue for the need for human-centered explainability to understand decisions from complex AI-enabled smart home applications.
- We also emphasize several human-computer interaction (HCI) methodologies, including user studies, prototyping, technology probes analysis, and heuristic evaluation, to achieve human-centered XAI-enabled smart home applications.

The rest of the paper is organized as follows: In Section 2, we present state-of-the-art technical XAI methods and human-centered XAI studies. In Section 3, we present two smart home application scenarios that illustrate the need for human-centered explainability in smart homes. We conducted experiments applying prominent ML techniques and XAI methods to predict and explain the models’ decisions in Section 4. We also demonstrate why the provided explanations are insufficient for users to understand them. In Section 5, we highlight multiple HCI methodologies that demonstrate how to elicit requirements and design human-centered explanations. Finally, Section 6 concludes the paper by outlining future directions.

## 2 Human-Centered XAI and Current Progress

The broad goal of XAI is to enable general users to understand the working principles of AI models and their decisions through explanations. Several terms, such as “*interpretable AI*” and “*transparent AI*”, are used interchangeably to describe the exact purpose of XAI [31]. The major objectives are similar: to make AI models and their decision-making processes understandable to general users through explanations. However, the number of methods focusing on a user-centered perspective is significantly lower than methods prioritizing improving model performance with technical explanations. As a result, the broad goal of XAI has not yet been fully achieved, which could potentially hinder the adoption of AI models in real-world applications.

This section provides a concise overview of the current cutting-edge approaches in explaining opaque decisions made by machine learning and deep learning-based predictive technologies, focusing on both technical and human-centered XAI. The primary aim of XAI advancements is to clarify the overall priorities and predictions of models for developers, facilitating the debugging process and enhancing model performance. While the field of technical XAI offers a wide range of techniques aimed at model improvement, it is important to note that the consideration of human-centered perspectives is relatively limited [31]. To present a comprehensive understanding of the literature, this section focuses on three key issues: i) technical XAI, which provides a background of XAI; ii)

human-centered XAI; iii) research gaps in the development of human-centered applications for smart home technology.

## 2.1 Technical XAI

The methods used in explainable artificial intelligence can be broadly classified into two categories: global and local explanation methods [33]. Global explainability methods aim to identify the overall priorities of a predictive model and provide a summary of the decision-making process. In contrast, local explainability methods focus on understanding why a specific predictive decision was made, shedding light on the insights associated with that decision. Additionally, explainable AI approaches can be categorized as either model-agnostic or model-specific. Model-agnostic approaches can be applied to any predictive model to explain its predictions, while model-specific explainable techniques are designed for particular predictive algorithms [19].

One prominent and widely used XAI method is SHapley Additive Explanation (SHAP) [34], which generally explains the global priorities of models and highlights the most and least significant features according to their contribution. Following the game theory concept, SHAP computes each feature’s weight that contributes to a specific decision. Moreover, SHAP can also provide local explanations for specific decisions. To explain deep neural network-based predictive models, well-known methods such as Grad-CAM [35], based on gradient localization, and Layer-wise Relevance Propagation (LRP) [36], which redistributes the output weight using backward propagation, can be used. These methods rely on saliency maps to explain decisions and can be applied to image and text-based applications. To uncover a DNN model’s decisions, DeepLIFT (Deep Learning Important FeaTures) [37] has been introduced to identify important features for specific predictions of a model. Lakkaraju et al. [38] proposed an XAI technique for AI experts to make the model and its behavior understandable by using subspace explanations. To provide explanations for experts, Schetinin et al. [39] introduced probabilistic interpretation for Bayesian decision tree models.

Another way to explain models’ predictions to users is through example-based explanations. With these types of explanations, XAI approaches provide similar instances that match the corresponding samples with the same decision. Several example-based interpretability models that consider similar prototypes and criticisms have been introduced to help users understand why a certain decision was made [40, 41]. To explain predictions for any complex deep learning model, Local Interpretable Model-Agnostic Explanations (LIME) [42] can create a *surrogate model* that mimics the performance of the complex model. The *surrogate model* is explainable, and its performance is quite similar to that of the original model. With LIME, any particular prediction can be explained for tabular and textual data by highlighting positive and negative features corresponding to the predicted decision.

To comprehend specific events, humans sometimes look for explanations that involve significant changes in the attributes of a sample, which can overturn the original decision. These are known as counterfactual explanations. They help

users understand “*why a different prediction was not possible?*” or “what changes could modify the final prediction?” Counterfactual explanations can also be illustrated by introducing new example samples that could reverse the decision, known as “what-if” scenarios. Various XAI approaches have been developed to explain “what-if” and counterfactual scenarios [43–45]. Other methods, such as the partial dependence plot (PDP), individual conditional expectation (ICE), and DiCE [44], are also used for similar purposes. To integrate multiple types of explanations into a unified framework, several open-source toolkits are available, including captum [46], AIX360 [47], and Anchor [48].

## 2.2 Human-centered XAI

The progress in the field of XAI, thus far, in generating technical explanations to interpret models for AI practitioners for debugging and improvement, far exceeds the advances made in human-centered perspectives of XAI. However, a significant number of scientific studies have emerged that focus on the user-centric perspective. In this section, we review some notable works on the evaluation of human-centered XAI.

Bell et al. [32] demonstrated that the model-agnostic explanations provided by one of the most prominent XAI techniques, SHAP, are not sufficiently comprehensible for general users. They assessed the effectiveness of these explanations through an empirical study involving non-technical participants in two distinct application areas, education, and finance [32]. Similarly, Abdul et al. [49] explored the trade-off between accuracy and simplicity in explanation presentation and proposed a cognitive generalized additive model (COGAM) for human-centered explanation delivery. A generalized XAI design principle was introduced to facilitate the presentation of local explanations to non-expert users by contextualizing the exploration of feature importance. An empirical study involving more than 80 participants was conducted to evaluate the effectiveness and user satisfaction with the explanations provided.

Chromik et al. [50] investigated the capability of non-expert users to understand and form a mental model of global explanations to comprehend the behavior of a model. Their findings indicated that global explanations are insufficient for understanding the overall model behavior. In a separate study, Hase et al. [51] measured the *simulatability* of various XAI techniques, including LIME, Anchor, prototypes, and decision boundaries, for tabular and textual data. They concluded that LIME achieved good simulatability in a few cases, and prototype models were useful for counterfactual explanations. Another study on simulatability found that current interpretability methods are inadequate in explaining model behavior [31, 52]. However, multiple studies assessing the effectiveness of incorporating XAI approaches in real-world decision-making concluded that the evidence does not align with the goals of XAI [53–56].

### 2.3 The Research Gap

The related works discussed above, addressing both technical and human-centered XAI, reveal a significant research gap in designing XAI systems for general users. Although current XAI systems can provide interpretable explanations in various forms—including global, local, counterfactual, and example-based explanations—these methods often fail to make predictions and model behavior understandable to non-expert users. Yet, one of the primary goals of XAI is to enable lay users to comprehend the predictions. Furthermore, applications in smart homes, such as energy demand forecasting, consumption practices, smart heating systems for comfortable home environments, and thermal comfort preferences with AI-enabled prediction systems, necessitate diverse perspectives for implementing user-centered XAI adoption.

Consider the complex task of developing a human-centered XAI energy prediction system for a smart home. This task is challenging due to the system’s internal complexity and the diverse backgrounds of non-expert users. While significant progress has been made in technical XAI for applications related to classification and regression, the complexity of energy forecasting, which involves two different dimensions—time and characteristics or features—makes it difficult to represent the underlying factors behind the forecast. Moreover, the explanations must be designed and presented in a human-centered manner, which poses additional challenges for creating human-centered explanations in smart home applications.

Another smart home application involves the automatic control of the heating system, which can be based on a prediction model of the inhabitants’ thermal preferences. This system generally utilizes physiological and environmental data from the inhabitants, as well as weather data, to predict their thermal comfort preferences [4, 5]. Based on these predictions, the smart home application then adjusts the heating system to control the indoor temperature. Since this system utilizes data from various sources, it is crucial for users to understand how and in what context their data are used. Therefore, we require human-centered, easily understandable explanations from such systems to ensure transparency and trust.

## 3 Human-Centered Explainability in Smart Home

Smart homes equipped with AI-driven applications, including energy demand forecasting, consumption routines analysis, heating systems monitoring, and controlling indoor temperature based on occupants’ thermal preference prediction systems, are becoming increasingly popular [1, 57]. These systems often use complex ML-based prediction systems that should be understandable to the household inhabitants [1, 5, 8]. The adoption of advanced XAI within complex smart home systems, offering human-centered explanations, could be transformative. As a result, non-expert smart home users would gain a clear and concise understanding of how their systems predict consumption and preferences, analyze their data for daily consumption practices, and control the indoor environment

to ensure a comfortable living space. Incorporating current XAI techniques and including users in the development loop to consider their perspectives could significantly enhance the provision of human-centered explanations.

We present an overview of human-centered XAI in smart home applications with occupants actively involved, as shown in Fig. 1. The prediction system should consider the occupants’ preferences and intentions to provide meaningful explanations. Initially, ML-enabled predictive models learn from preprocessed data to make future predictions and forecasts. When implementing XAI models to elucidate the decision-making processes of these models and offer explanations for individual predictions, it is crucial to incorporate a human-centered perspective. In this context, human-computer interaction methods are essential for analyzing user feedback. Consequently, effectively presenting these explanations can enhance the adoption of AI-enabled systems in real-world smart home settings. To highlight the importance of human-centered explainability methods in smart home applications, we selected two relevant and extensively investigated problems within the smart home context.

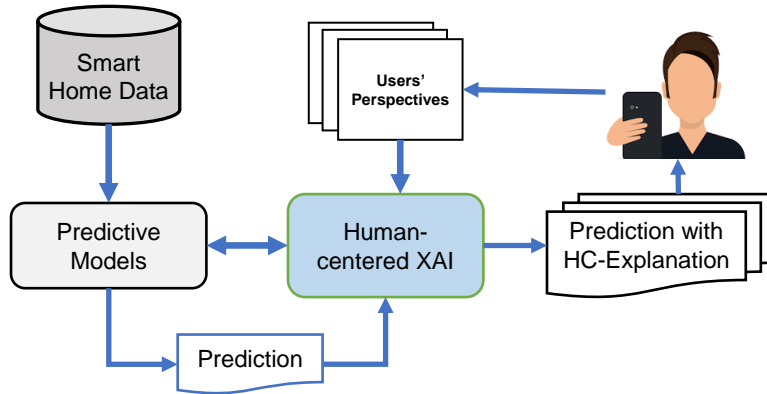


Fig. 1: An overview of a human-centered XAI-enabled Smart Home systems

To select the problem domain in the context of smart homes, we focus on applications related to household energy consumption. This emphasis allows us to explore areas that enhance resident comfort while also optimizing energy use. Additionally, smart applications should offer insights to help people optimize and reduce their energy costs. Therefore, we have chosen two applications: household energy demand forecasting and automatic control of heating based on predictions of thermal comfort preferences. Both applications are concerned with energy consumption practices. The first application is an energy demand forecasting system designed to increase awareness and optimize household electricity use. The second involves modeling occupants’ thermal comfort preferences to ensure a comfortable and healthy indoor environment. In the remainder of



this section, we present two relevant smart home sub-tasks that require human-centered explanations for successful adoption.

### 3.1 Household energy demand forecasting

Energy consumption in residential and commercial buildings significantly exceeds that in other sectors [1, 5]. Additionally, the price of energy is continuously increasing worldwide. However, this heightened consumption also results in significant CO2 emissions, posing a serious threat to global warming and the environment [5]. Smart home systems address these concerns by providing future energy demand forecasts for households based on historical energy usage data collected from various appliances, thanks to advanced sensor technology [8]. Such predictions of total energy demand for the upcoming month or week might raise household members’ awareness of their energy-related activities, potentially encouraging them to optimize their energy usage.

Nevertheless, energy demand forecasting systems typically rely on highly sophisticated ML and DL techniques [8, 9, 24], which perform complex calculations and lead to opaque decision-making processes. Consequently, some predictions may surprise users with unexpected outcomes. For example, if a forecasting system predicts a high (or low) total energy consumption for washing machines in the next month, users might be taken aback, as they might perceive washing machines to consume less (or more) energy than predicted. In such scenarios, users require easily understandable explanations from the systems to comprehend and build trust in AI-enabled systems, thereby enhancing their adoption.

Unlike other classification and regression tasks, providing explanations for energy demand forecasting systems is non-trivial, as it involves multivariate time series forecasting [1]. The explanations in such systems cover two dimensions: they relate to features or attributes and time. Consequently, explainability methods must capture the impact of time and features, making understanding explanations for time series forecasting more challenging for users. Research into human-centered explainability methods in this field is crucial to address this issue. This research will enable inhabitants to understand why a certain energy demand is anticipated for the upcoming month and will facilitate the development of more optimal energy consumption plans based on factual explanations.

### 3.2 Occupants’ thermal comfort preference modeling

Indoor thermal comfort is essential for the well-being, comfort, and work productivity of inhabitants [4, 5]. With recent advancements in efficient sensors and smart home appliances, AI-driven heating, ventilation, and air conditioning (HVAC) systems can monitor and control the indoor environment. These systems have gained considerable attention for applying machine learning techniques to automatically control parameters related to the comfort of the indoor environment [4, 58]. Personal thermal comfort preferences vary widely from person to person, making the prediction of individual preferences crucial for providing occupant-level comfort in households [4, 5]. Based on these preference

predictions, the heating system can be automatically adjusted to control the temperature at the occupant’s location.

Householders often struggle with inconsistent temperatures in their homes, especially during extreme weather conditions, leading to discomfort and high energy costs [59]. Maintaining a consistent temperature throughout the house can be challenging, causing the HVAC system to work harder to maintain a comfortable temperature using AI-enabled computational models. However, these complex AI systems often lack interpretability. Traditional temperature control systems typically react by adjusting the temperature only after it has already started to fluctuate, resulting in uncomfortable temperature swings and inefficiencies. To address this problem, a more proactive approach to temperature control is needed. Various methods have been introduced to predict personal thermal comfort preferences using complex machine learning and deep learning models [5, 11, 15, 60]. These automated systems, powered by complex AI models, can monitor and control the indoor environment. However, the decision-making process and the rationale behind specific predictions and actions often remain unclear to the inhabitants, including AI practitioners themselves. As XAI methods progress, it is crucial to make the explanations human-centered, enabling household occupants to understand the reasons behind predictions and the decision-making of the models. This would facilitate the adoption of such complex models and ensure successful smart home applications.

## 4 Experiments and Analysis

This section presents the details of the experiments we conducted on the two aforementioned smart home scenarios. We carried out experiments by training various predictive models on two distinct datasets collected for both applications. Initially, we trained predictive models and subsequently applied two well-known XAI methods, namely SHAP [34] and DeepLIFT [37]. We then analyzed the generated explanations for both smart home applications and sought to identify reasons why these explanations might not be sufficient for smart home users to comprehend. This section presents the experimental results, generated explanations, and their analysis for each scenario.

### 4.1 Energy demand forecasting in smart home

**Experimental Settings:** We conducted experiments on the REFIT dataset [61], which includes energy consumption data collected from 20 diverse households. The data encompasses various appliances such as the Fridge-Freezer, Tumble Dryer, Washing Machine, Dishwasher, Desktop Computer, Television, Microwave, Kettle, and Toaster, with energy consumption recorded at 8-second intervals. We modeled the weekly energy demand forecasting problem using a classical LSTM-based model. The LSTM-based forecasting model features 10 total features, a sequence length of 7, two hidden layers, 64 hidden units in each layer, 100 epochs,

Table 1: Prediction performance in forecasting energy demand on 4 different households of REFIT dataset

House	MAE	MAPE	MSE	RMSE
House 2	0.1351	0.5421	0.03	0.1752
House 5	0.0768	0.4075	0.01	0.10
House 8	0.1934	0.4618	0.0451	0.2095
House 13	0.0435	0.8175	0.003	0.0564

a learning rate of 0.001, and a batch size of 64.

**Results:** The performance of the LSTM-based forecasting model to predict upcoming weekly energy demand is presented in Table 1. We can see that the performance is quite effective in terms of four different evaluation metrics, including mean squared error (MSE), Root-MSE (RMSE), mean average error (MAE), and Mean absolute percentage error (MAPE). The performance across different households varied widely. For house 5, the forecasting performance is better than that of another household in terms of MAPE and MSE. On the other hand, for house 13, LSTM achieved the best performance in terms of MAE and RMSE. However, presenting the forecasting performance here makes sense in that the generated explanations for decisions can better capture facts and reasons.

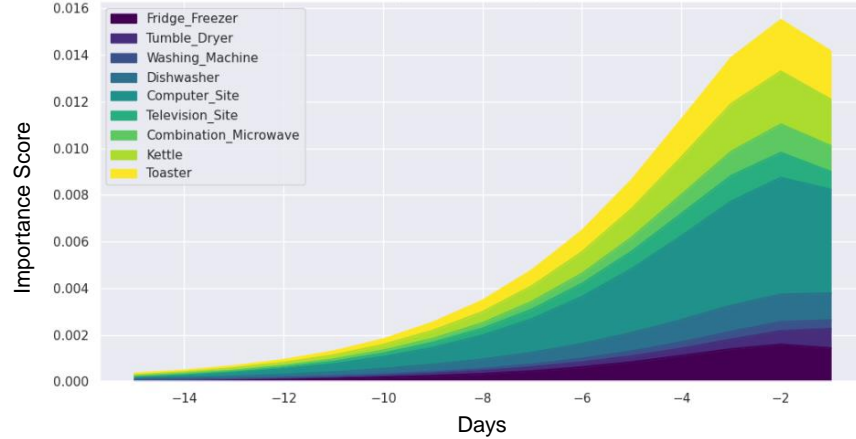


Fig. 2: Explanations for weekly energy demand forecasting highlighting the contributions of different appliances.

**Explainability:** To explain the predictions from the model, we utilized Deep Learning Important Feature (DeepLIFT), which approximates Shapley values to provide an explanation. This method combines the contributions of different appliances towards the overall prediction, thereby informing users about the activities responsible for the total energy consumption. The explanations for

weekly energy forecasting, presented in Figure 2, show how the contributions of different appliances vary over time. These explanations are fairly technical, indicating that the contributions of different appliances change with time. Previous studies have shown that general users often struggle to understand even straightforward explanations generated for binary classification tasks.

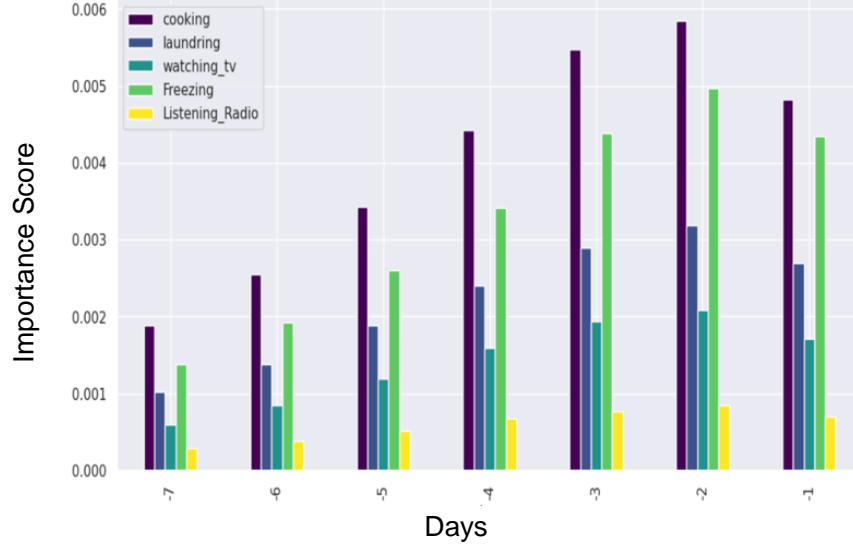


Fig. 3: Explanations for weekly energy demand forecasting highlighting consumption activity corresponding to the time (day)

On the other hand, energy demand forecasting is an even more challenging task, and the explanations provided differ significantly. The dimension of these explanations also relates to time. As a result, there is a need for explanations that help laypeople understand how these AI-enabled forecasting systems make decisions in their homes. To simplify, we sum up the contributions of different appliances and present the explanation using a bar chart. Figure 3 shows the contributions of different household activities responsible for overall energy consumption, with cooking activities expected to consume the most energy. What use are these explanations to a user? While they provide some insight into the activities most responsible for energy use, they do not enable users to optimize their energy consumption practices with the current form of explanations. It is also clear that understanding explanations generated by a single successful XAI technique can be challenging. The impact of different activities on overall consumption varies across different days. Non-expert users may develop a negative perception when confronted with such complex explanations. Therefore, explanations should be easy to follow, trustworthy, and tailored to a human-centered perspective to facilitate sense-making.

Table 2: The performance of different classical machine learning models on predicting personal thermal comfort preference. The best results are in **bold**.

Model	Kappa	Accuracy	AUC
Decision Tree	0.6609	0.8315	0.8810
Support Vector Machine	0.5470	0.8167	0.9198
K-nearest Neighbor	0.3810	0.7589	0.8101
Gaussian Naive Bayes	0.4311	0.7148	0.7689
XGBoost	<b>0.6774</b>	<b>0.8457</b>	<b>0.9487</b>
Random Forest	0.5901	0.8195	0.9258

## 4.2 Personal thermal comfort preference prediction

**Experimental Settings:** For the second smart home sub-task, “*thermal comfort preference modeling*,” we conducted experiments using a wearable dataset collected by UC Berkeley. The dataset originates from a field experiment involving 14 subjects living in Berkeley and San Francisco [4]. It contains a total of 3848 samples from all participants, categorized into their thermal comfort preferences: “Cooler” (class 0), “No Change” (class 1), and “Warmer” (class 2). Further details about the dataset can be found in [4]. After preprocessing the values of the features, we applied a feature selection technique to identify relevant features, resulting in the selection of 32 different features [4]. We then trained six different prominent classical machine learning models, including Decision Tree (DT), Support Vector Machine (SVM), K-nearest Neighbor (KNN), Gaussian Naive Bayes (GNB), XGBoost (XGB), and Random Forest (RF).

**Results:** The performance of all trained machine learning models is presented in Table 2. We evaluated their performance using metrics such as Cohen’s Kappa, Accuracy, and Area Under the Curve (AUC). Given that the dataset was quite imbalanced in terms of the number of samples for different classes, we employed metrics that can effectively evaluate the performance of ML classifiers on imbalanced data. For this purpose, we applied Cohen’s Kappa and AUC. From Table 2, it is evident that XGBoost achieved better performance across all evaluation metrics. The performance of the other five models was also quite consistent compared to the baseline models on the same dataset by [4].

**Explainability:** We incorporated the prominent XAI method SHAP to generate explanations. We illustrated the global and local explanations in Figure 4 and 5. Figure 4 shows the global explanations, indicating that thermal sensitivity is the most important feature for modeling personal thermal comfort. The next most important features are cold experience, age, weight, and work hours. As an AI practitioner, one can understand which features the model prioritizes for the overall decision. However, as an inhabitant of the household, some might struggle to derive meaningful insights from this explanation.

In Fig. 5, we illustrate the explanations for a particular subject’s thermal preference at a certain time in terms of a waterfall plot; these explanations fol-




Fig. 4: Global explanation for personal thermal comfort preference prediction highlighting model’s overall priorities.

low the global explanations in a broader sense. We can observe that the most important features are thermal sensitivity, height, and temperature in the ankle. Once again, developers can leverage these insights to enhance the models’ performance by canonicalizing and modifying them. Nevertheless, general occupants may find it challenging to derive actionable insights from this information.

## 5 Challenges and HCI Techniques for Human-centered Explainability

After carefully analyzing the current literature on XAI and human-centered XAI (Section 2) and generating explanations using two established XAI methods across two different smart home scenarios, we have identified challenges that need addressing to achieve human-centered XAI in these applications. Smart home applications span various problem domains, including classification (e.g., PTC preference prediction), time-series forecasting (e.g., energy demand forecasting, energy billing), and regression (e.g., activity recognition). Thus, the challenges for achieving user-centric explainability must be tailored to each specific domain. While technical explanations are useful for debugging and optimizing

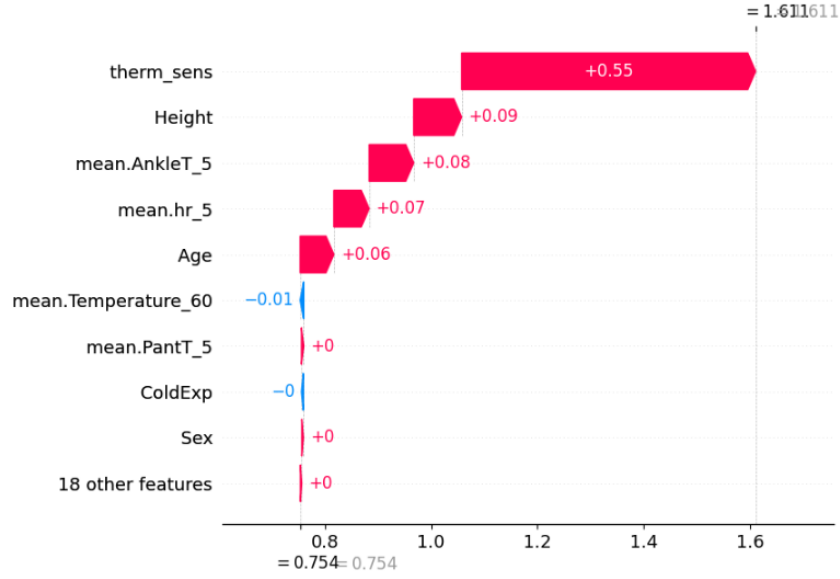


Fig. 5: Explanation for a decision that predicts the occupants felling “warmer”.

model performance, they do not typically aid general users’ decision-making. If explanations do not meet users’ needs, this can result in a lack of trust in the model and resistance to its use [22]. Human-centered XAI tools, designed with end-users in mind and providing contextually relevant and understandable explanations, can overcome these challenges [23]. To offer human-centered explanations for smart home systems, it is crucial to employ user-focused XAI tools, provide understandable explanations, and address the users’ inquiries and concerns [23]. Addressing the following challenges will help in generating human-centered, XAI-enabled explanations, ultimately improving the adoption of AI systems in real-world applications.

### 5.1 Challenges in achieving human-centered explainability

When discussing the challenges of human-centered explainability, we adopt the notions of syntax, semantics, and pragmatics, which are traditionally utilized in linguistics but are effectively applicable in identifying challenges in smart home explanations. These concepts aid in understanding and improving how users design and perceive explanations.

**The syntax level of explanations** refers to the presentation and visual encoding of explanations, emphasizing user-friendly choices in colors, fonts, layouts, and chart design to minimize cognitive load and simplify interactions with smart devices [62]. Although charts generated by XAI frameworks such as SHAP or LIME are highly accurate, they tend to be mathematical and unengaging. In technical domains, this might not pose a problem; however, for lay smart home users, factors such as aesthetics, playfulness, comprehensibility, appropriate mea-

surement units, and careful wording are critical [63]. For instance, Schwartz et al. [64] demonstrate that technical units like *kWh* (kilowatt-hour) or *kg CO<sub>2</sub> eq* (equivalent to the effect of one kg of CO<sub>2</sub> emission) are too complex, whereas laypersons typically prefer money as a well-known, easy-to-interpret unit. Another challenge in the visual design of explanations is the small screen size [62], as most people interact with their smart home through smartphones or wall-mounted interfaces, necessitating the simplification of complex explanations for small-screen visualization.

**The semantic level of explanations** concerns how explanations are interpreted and what mental models are generated [62]. In the context of smart homes, for instance, many individuals possess incorrect mental models of heating systems, leading to improper heating behaviors [65]. Therefore, explanations must be mathematically precise and assist users in developing accurate mental models. In this regard, Schwartz et al. [64] demonstrate that people often rely on ethno- or folk-methods to construct their mental models. Regarding domestic energy consumption, people interpret the information provided by eco-feedback systems using money as the preferred unit to assess appliance consumption, relate consumption to their habits, or compare their consumption with others' [64]. Explanations should leverage these folk methods to help people build accurate mental models [64, 65].

**The pragmatic level of explanations** refers to the context-dependent, practical significance of explanations as they apply to the user's daily life. For example, explanations can serve various purposes, such as building trust in smart home systems [66], increasing the energy literacy of residents [67], supporting reflection on wasteful consumption habits [64], and prompting actions to detect and replace inefficient appliances. Additionally, they enhance predictability and accountability [66] and support the co-performance of controlling domestic appliances [68].

Explanations should be tailored to these pragmatic considerations to be effective. For example, simple recommendations accompanied by "what-if" explanations [45] are more effective for actions such as detecting and replacing wasteful appliances. In contrast, more elaborated, cause-effect-oriented explanation approaches [69] are more helpful in enhancing energy literacy. By considering all three aspects—syntax, semantics, and pragmatics—we can better identify and overcome the specific challenges associated with explaining smart home technologies, making them more user-friendly and aligned with human-centered design principles.

In the following, we discuss these challenges across three scenarios.

#### 1. **Making predictions and autonomous actions interpretable for users:**

In the context of energy demand forecasting, XAI tools can provide understandable explanations for the predictions made by the forecasting model, highlighting the most influential factors contributing to the prediction. These explanations should be presented in a user-friendly format that aligns with the user's intuition and understanding of the problem. Visualizations such as graphs and charts can be used to illustrate the data and make it easier for



users to comprehend the predictions. In the case of an autonomous action performed by an AI-enabled system (e.g., adjusting the room temperature based on predicted thermal comfort preferences), the system should provide clear explanations for why that particular action was taken.

2. **Providing insights for an optimal energy consumption plan:** Human-centered, XAI-powered smart home applications should offer users insights into their optimal energy consumption plan by considering their energy usage patterns, environmental conditions, and preferences. The system might involve developing a personalized energy optimization plan, such as adjusting the temperature based on the user’s daily routines or turning off lights in unoccupied rooms. The human-centered XAI-enabled system should provide suggestions and explanations for energy-saving practices and offer feedback on the impact of these practices on energy consumption.
3. **Making users aware of energy consumption:** XAI tools should help users understand how their energy consumption patterns affect the home environment and their energy bills. This can be achieved by providing real-time feedback on energy usage, highlighting areas where energy could be conserved, and suggesting energy-efficient practices. For instance, the XAI tool can send alerts or notifications to remind users to turn off lights or appliances when not in use.




Fig. 6: HCI techniques to enhance human-centered explainability

## 5.2 HCI techniques to enhance human-centered explainability

Focusing on human-centered XAI can align with users’ needs and preferences, increasing their trust and understanding of smart home AI systems [23]. This alignment is particularly significant as we strive for a more energy-efficient and comfortable living environment with smart home systems. Thus, Incorporating HCI methodologies becomes effective in advancing smart home, human-centered XAI implementation. The following non-linear process in Fig. 6 illustrates this approach: User Studies [70], Prototyping [71], Technology Probes Analysis [72], and Heuristic Evaluation [73].

1. **User Studies:** User studies employ detailed observation, daily tracking, and interviews based on praxeological grounded design methods to understand contextual user practices in smart home environments [70]. This approach

aims to gather deep insights into how users interact with XAI artifacts of smart home technology, facilitating a refined understanding of social practices. By analyzing these interactions, developers can align XAI systems more closely with user habits and expectations. Such alignment increases user trust and enhances system transparency. These insights are crucial for tailoring AI functionalities to manage home heating or cooling systems efficiently. For example, based on collected data, AI could predict and adjust indoor temperatures to optimal levels just before users return home or during specific weather conditions, thus maximizing comfort without the need for manual adjustments [1, 4, 5]. These adaptive adjustments contribute significantly to improving energy efficiency and overall user comfort.

2. **Prototyping:** Developing several low-fidelity prototypes offers a cost-effective way to explore different designs and continuously gather user feedback. This process is characterized by iterative testing and refinement cycles, heavily using participatory design principles to align with user expectations and improve usability [71]. Through prototyping, designers can quickly adapt and evolve XAI features based on real user feedback, ensuring the system is both intuitive and directly useful to end-users. Such an approach allows developers to fine-tune how XAI communicates energy-saving decisions to users. For example, the explanations might suggest reducing electricity use during off-peak hours as a cost-saving measure. By clearly explaining the rationale behind such recommendations, prototyping enhances the usability of XAI, making it easier for users to trust and follow the AI's guidance. This process not only aids in reducing energy costs but also helps in educating users about efficient energy practices.
3. **Technology Probes Analysis:** Technology probes are invaluable for understanding how users interact with and respond to XAI systems within smart homes, especially focusing on aspects such as interpretability, responsibility, and relevance of AI explanations. By deploying technology probes, such as smart meters that track energy usage and provide feedback and advice based on AI analysis, developers can gather rich data on user behavior and preferences [72]. These probes reveal user reactions to automated suggestions for optimizing energy consumption. The insights gained from technology probes allow developers to refine human-centered XAI explanations, ensuring they are both meaningful and actionable. This process enhances the system's usability and boosts users' understanding of and trust in the AI explanations, leading to more effective and sustainable smart energy management.
4. **Heuristic Evaluation:** Heuristic evaluation involves collaboration with subject matter experts, including HVAC, electrical, and social engineers, to assess AI systems' responsibility and user-centric design. This evaluation focuses on key aspects such as transparency, user control, and ethical considerations. Experts examine whether an AI system's explanations for recommending specific energy-saving measures are understandable and align with user values [73]. Such evaluations are crucial because they help ensure that AI systems are designed with a strong emphasis on user-centric princi-

ples, which promotes a better understanding of and trust in the technology. Heuristic evaluation is particularly effective at identifying aspects of AI explanations that may not be evident through end-user testing alone. Doing so aids users in making informed decisions about their energy use, like understanding why certain settings are recommended for maximizing thermal comfort without excessive energy use [1, 4, 5].

## 6 Conclusions and Future Directions

Our research strives to achieve human-centered XAI for smart home applications, aiming to make complex AI-driven models understandable to laypeople. Through experiments on two smart home sub-tasks, we demonstrated the challenges associated with understanding decisions from such applications. We argue that current explanation generation techniques are insufficient for making general users comprehend these decisions. By identifying major challenges, we highlight the need for human-centered explainability and discuss how these challenges can be addressed using various human-computer interaction (HCI) techniques, including user feedback, co-performance considerations, and expert-user co-design. Future human-centered XAI can apply HCI techniques to understand users' requirements and preferences better, enabling them to understand decisions from AI-driven systems. Based on the outcomes of these techniques, we aim to develop human-centered explanations that facilitate user understanding and action.

However, challenges persist in generating human-centered explanations that foster trust and interpretability in AI systems. Future research directions include:

- Exploring natural language explanations and interactive interfaces
- Developing standardized frameworks to evaluate the general user experience of human-centered explainability
- Integrating HCI methodologies into the XAI development life-cycle to ensure user-centered and effective explanations

Future research on emerging user interfaces, including ubiquitous and pervasive technologies, can also advance the current state of the art on human-centered XAI in the context of the smart home energy domain. As we envision the integration of smart devices such as voice assistants, smart watches, and embedded sensors into everyday environments, these technologies provide a rich platform for deploying human-centered solutions to benefit sustainability. Such interfaces can offer intuitive and context-aware interactions, making AI explanations part of the natural user environment and activities.

## Acknowledgment

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 955422.

## References

1. Md Shajalal, Alexander Boden, and Gunnar Stevens. Towards user-centered explainable energy demand forecasting systems. In *Proceedings of the Thirteenth ACM International Conference on Future Energy Systems*, pages 446–447, 2022.
2. Jin-Young Kim and Sung-Bae Cho. Electric energy demand forecasting with explainable time-series modeling. In *2020 International Conference on Data Mining Workshops (ICDMW)*, pages 711–716. IEEE, 2020.
3. Alper T Alan, Mike Shann, Enrico Costanza, Sarvapali D Ramchurn, and Sven Seuken. It is too hot: An in-situ study of three designs for heating. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 5262–5273, 2016.
4. Md Shajalal, Milad Bohlouli, Hari Prasanna Das, Alexander Boden, and Gunnar Stevens. Improved thermal comfort model leveraging conditional tabular gan focusing on feature selection. *IEEE Access*, 12:30039–30053, 2024.
5. Md Shajalal, Milad Bohlouli, Hari Prasanna Das, Alexander Boden, and Gunnar Stevens. Focus on what matters: improved feature selection techniques for personal thermal comfort modelling. In *Proceedings of the 9th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pages 496–499, 2022.
6. Lina Stankovic, Vladimir Stankovic, Jing Liao, and Clevo Wilson. Measuring the energy intensity of domestic activities from smart meter data. *Applied Energy*, 183:1565–1580, 2016.
7. Haider Mshali, Tayeb Lemlouma, Maria Moloney, and Damien Magoni. A survey on health monitoring systems for health smart homes. *International Journal of Industrial Ergonomics*, 66:26–56, 2018.
8. Jin-Young Kim and Sung-Bae Cho. Explainable prediction of electric energy demand using a deep autoencoder with interpretable latent space. *Expert Systems with Applications*, 186:115842, 2021.
9. Jin-Young Kim and Sung-Bae Cho. Electric energy consumption prediction by deep learning with state explainable autoencoder. *Energies*, 12(4):739, 2019.
10. Nicolai Bo Vanting, Zheng Ma, and Bo Nørregaard Jørgensen. A scoping review of deep neural networks for electric load forecasting. *Energy Informatics*, 4(2):1–13, 2021.
11. Mahmoud M Abdelrahman, Adrian Chong, and Clayton Miller. Personal thermal comfort models using digital twins: Preference prediction with bim-extracted spatial-temporal proximity data from build2vec. *Building and Environment*, 207:108532, 2022.
12. Aniruddh Chennapragada, Divya Periyakoil, Hari Prasanna Das, and Costas J Spanos. Time series-based deep learning model for personal thermal comfort prediction. In *Proceedings of the Thirteenth ACM International Conference on Future Energy Systems*, pages 552–555, 2022.
13. Nasim Eslamirad, Soheil Malekpour Kolbadinejad, Mohammadjavad Mahdavinnejad, and Mohammad Mehranrad. Thermal comfort prediction by applying supervised machine learning in green sidewalks of tehran. *Smart and Sustainable Built Environment*, 9(4):361–374, 2020.
14. Nan Gao, Wei Shao, Mohammad Saiedur Rahaman, Jun Zhai, Klaus David, and Flora D Salim. Transfer learning for thermal comfort prediction in multiple cities. *Building and Environment*, 195:107725, 2021.

15. Nivethitha Somu, Anirudh Sriram, Anupama Kowli, and Krithi Ramamritham. A hybrid deep transfer learning strategy for thermal comfort prediction in buildings. *Building and Environment*, 204:108133, 2021.
16. Nikos D Sakkas, Sofia Yfanti, Pooja Shah, Nikitas Sakkas, Christina Chaniotakis, Costas Daskalakis, Eduard Barbu, and Marharyta Domnich. Explainable approaches for forecasting building electricity consumption. 2023.
17. Matias Quintana, Stefano Schiavon, Kwok Wai Tham, and Clayton Miller. Balancing thermal comfort datasets: We gan, but should we? In *Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pages 120–129, 2020.
18. Md Shajalal, Sebastian Deneff, Md Rezaul Karim, Alexander Boden, and Gunnar Stevens. Unveiling black-boxes: Explainable deep learning models for patent classification. In *World Conference on Explainable Artificial Intelligence*, pages 457–474. Springer, 2023.
19. Md Shajalal, Alexander Boden, and Gunnar Stevens. Explainable product backorder prediction exploiting cnn: Introducing explainable models in businesses. *Electronic Markets*, pages 1–16, 2022.
20. Md Rezaul Karim, Md Shajalal, Alexander Graß, Till Döhmen, Sisay Adugna Chala, Alexander Boden, Christian Beecks, and Stefan Decker. Interpreting black-box machine learning models for high dimensional datasets. In *2023 IEEE 10th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 1–10. IEEE, 2023.
21. Md Rezaul Karim, Tanhim Islam, Md Shajalal, Oya Beyan, Christoph Lange, Michael Cochez, Dietrich Rebholz-Schuhmann, and Stefan Decker. Explainable ai for bioinformatics: Methods, tools and applications. *Briefings in bioinformatics*, 24(5):bbad236, 2023.
22. Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Benoitot, Siham Tabik, Alberto Barbado, Salvador García, Sergio Gil-López, Daniel Molina, Richard Benjamins, et al. Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai. *Information fusion*, 58:82–115, 2020.
23. Upol Ehsan, Philipp Wintersberger, Q Vera Liao, Elizabeth Anne Watkins, Carina Manger, Hal Daumé III, Andreas Riener, and Mark O Riedl. Human-centered explainable ai (hcxai): beyond opening the black-box of ai. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, pages 1–7, 2022.
24. Upol Ehsan, Philipp Wintersberger, Q Vera Liao, Martina Mara, Marc Streit, Sandra Wachter, Andreas Riener, and Mark O Riedl. Operationalizing human-centered perspectives in explainable ai. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–6, 2021.
25. M Humayn Kabir, Khondokar Fida Hasan, Mohammad Kamrul Hasan, and Keyvan Ansari. Explainable artificial intelligence for smart city application: A secure and trusted platform. *arXiv preprint arXiv:2111.00601*, 2021.
26. Henrik Mucha, Sebastian Robert, Rüdiger Breitschwerdt, and Michael Fellmann. Towards participatory design spaces for explainable ai interfaces in expert domains. In *43rd German Conference on Artificial Intelligence, Bamberg, Germany*, 2020.
27. Ana I Grimaldo and Jasminko Novak. Combining machine learning with visual analytics for explainable forecasting of energy demand in prosumer scenarios. *Procedia Computer Science*, 175:525–532, 2020.
28. Daniele Riboni. Keynote: Explainable ai in pervasive healthcare: Open challenges and research directions. In *2021 IEEE International Conference on Pervasive*

- Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)*, pages 1–1. IEEE, 2021.
29. Arun Rai. Explainable ai: From black box to glass box. *Journal of the Academy of Marketing Science*, 48(1):137–141, 2020.
  30. Nico Castelli, Corinna Ogonowski, Timo Jakobi, Martin Stein, Gunnar Stevens, and Volker Wulf. What happened in my home? an end-user development approach for smart home data visualization. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 853–866, 2017.
  31. Yao Rong, Tobias Leemann, Thai-trang Nguyen, Lisa Fiedler, Tina Seidel, Gjergji Kasneci, and Enkelejda Kasneci. Towards human-centered explainable ai: User studies for model explanations. *arXiv preprint arXiv:2210.11584*, 2022.
  32. Andrew Bell, Ian Solano-Kamaiko, Oded Nov, and Julia Stoyanovich. It’s just not that simple: an empirical study of the accuracy-explainability trade-off in machine learning for public policy. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 248–266, 2022.
  33. Christoph Molnar. *Interpretable machine learning*. Lulu. com, 2020.
  34. Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.
  35. Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
  36. Grégoire Montavon, Alexander Binder, Sebastian Lapuschkin, Wojciech Samek, and Klaus-Robert Müller. Layer-wise relevance propagation: an overview. *Explainable AI: interpreting, explaining and visualizing deep learning*, pages 193–209, 2019.
  37. Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. In *International conference on machine learning*, pages 3145–3153. PMLR, 2017.
  38. Himabindu Lakkaraju, Ece Kamar, Rich Caruana, and Jure Leskovec. Faithful and customizable explanations of black box models. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 131–138, 2019.
  39. Vitaly Schetin, Jonathan E Fieldsend, Derek Partridge, Timothy J Coats, Wojtek J Krzanowski, Richard M Everson, Trevor C Bailey, and Adolfo Hernandez. Confident interpretation of bayesian decision tree ensembles for clinical applications. *IEEE Transactions on Information Technology in Biomedicine*, 11(3):312–319, 2007.
  40. Been Kim, Rajiv Khanna, and Oluwasanmi O Koyejo. Examples are not enough, learn to criticize! criticism for interpretability. *Advances in neural information processing systems*, 29, 2016.
  41. Karthik S Gurumoorthy, Amit Dhurandhar, Guillermo Cecchi, and Charu Aggarwal. Efficient data representation by selecting prototypes with importance weights. In *2019 IEEE International Conference on Data Mining (ICDM)*, pages 260–269. IEEE, 2019.
  42. Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. ” why should i trust you?” explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144, 2016.
  43. Arnaud Van Looveren and Janis Klaise. Interpretable counterfactual explanations guided by prototypes. In *Machine Learning and Knowledge Discovery in Databases*.

- Research Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part II 21*, pages 650–665. Springer, 2021.
44. Ramaravind K Mothilal, Amit Sharma, and Chenhao Tan. Explaining machine learning classifiers through diverse counterfactual explanations. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 607–617, 2020.
  45. Sandra Wachter, Brent Mittelstadt, and Chris Russell. Counterfactual explanations without opening the black box: Automated decisions and the gdpr. *Harv. JL & Tech.*, 31:841, 2017.
  46. Narine Kokhlikyan, Vivek Miglani, Miguel Martin, Edward Wang, Bilal Alsallakh, Jonathan Reynolds, Alexander Melnikov, Natalia Kliushkina, Carlos Araya, Siqi Yan, et al. Captum: A unified and generic model interpretability library for pytorch. *arXiv preprint arXiv:2009.07896*, 2020.
  47. Vijay Arya, Rachel KE Bellamy, Pin-Yu Chen, Amit Dhurandhar, Michael Hind, Samuel C Hoffman, Stephanie Houde, Q Vera Liao, Ronny Luss, Aleksandra Mojsilović, et al. Ai explainability 360 toolkit. In *Proceedings of the 3rd ACM India Joint International Conference on Data Science & Management of Data (8th ACM IKDD CODS & 26th COMAD)*, pages 376–379, 2021.
  48. Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Anchors: High-precision model-agnostic explanations. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
  49. Ashraf Abdul, Christian von der Weth, Mohan Kankanhalli, and Brian Y Lim. Cogam: measuring and moderating cognitive load in machine learning model explanations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2020.
  50. Michael Chromik, Malin Eiband, Felicitas Buchner, Adrian Krüger, and Andreas Butz. I think i get your point, ai! the illusion of explanatory depth in explainable ai. In *26th International Conference on Intelligent User Interfaces*, pages 307–317, 2021.
  51. Peter Hase and Mohit Bansal. Evaluating explainable ai: Which algorithmic explanations help users predict model behavior? *arXiv preprint arXiv:2005.01831*, 2020.
  52. Zachary C Lipton. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3):31–57, 2018.
  53. Q Vera Liao and Kush R Varshney. Human-centered explainable ai (xai): From algorithms to user experiences. *arXiv preprint arXiv:2110.10790*, 2021.
  54. Gagan Bansal, Tongshuang Wu, Joyce Zhou, Raymond Fok, Besmira Nushi, Ece Kamar, Marco Tulio Ribeiro, and Daniel Weld. Does the whole exceed its parts? the effect of ai explanations on complementary team performance. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–16, 2021.
  55. Forough Poursabzi-Sangdeh, Daniel G Goldstein, Jake M Hofman, Jennifer Wortman Wortman Vaughan, and Hanna Wallach. Manipulating and measuring model interpretability. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, pages 1–52, 2021.
  56. Xinru Wang and Ming Yin. Effects of explanations in ai-assisted decision making: principles and comparisons. *ACM Transactions on Interactive Intelligent Systems*, 12(4):1–36, 2022.

57. Rikke Hagensby Jensen, Yolande Strengers, Jesper Kjeldskov, Larissa Nicholls, and Mikael B Skov. Designing the desirable smart home: A study of household experiences and energy consumption impacts. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2018.
58. Mike Shann, Alper Alan, Sven Seuken, Enrico Costanza, and Sarvapali D Ramchurn. Save money or feel cozy?: A field experiment evaluation of a smart thermostat that learns heating preferences. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, volume 16. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), 2017.
59. Véronique Vasseur, Anne-Francoise Marique, and Vladimir Udalov. A conceptual framework to understand households’ energy consumption. *Energies*, 12(22):4250, 2019.
60. Rocío Escandón, Fabrizio Ascione, Nicola Bianco, Gerardo Maria Mauro, Rafael Suárez, and Juan José Sendra. Thermal comfort prediction in a building category: Artificial neural network generation from calibrated models for a social housing stock in southern europe. *Applied Thermal Engineering*, 150:492–505, 2019.
61. David Murray, Lina Stankovic, and Vladimir Stankovic. An electrical load measurements dataset of united kingdom households from a two-year longitudinal study. *Scientific data*, 4(1):1–12, 2017.
62. Robert Spence. *Information visualization*, volume 1. Springer, 2001.
63. Nico Castelli, Gunnar Stevens, and Timo Jakobi. Information visualization at home: A literature survey of consumption feedback design. 2019.
64. Tobias Schwartz, Gunnar Stevens, Leonardo Ramirez, and Volker Wulf. Uncovering practices of making energy consumption accountable: A phenomenological inquiry. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 20(2):1–30, 2013.
65. Willett Kempton. Two theories of home heat control. *Cognitive science*, 10(1):75–90, 1986.
66. Timo Jakobi, Gunnar Stevens, Nico Castelli, Corinna Ogonowski, Florian Schaub, Nils Vindice, Dave Randall, Peter Tolmie, and Volker Wulf. Evolving needs in iot control and accountability: A longitudinal study on smart home intelligibility. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(4):1–28, 2018.
67. Tobias Schwartz, Sebastian Deneff, Gunnar Stevens, Leonardo Ramirez, and Volker Wulf. Cultivating energy literacy: results from a longitudinal living lab study of a home energy management system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1193–1202, 2013.
68. Lenneke Kuijter and Elisa Giaccardi. Co-performance: Conceptualizing the role of artificial agency in the design of everyday life. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2018.
69. Robert R Hoffman, Shane T Mueller, and Gary Klein. Explaining explanation, part 2: Empirical foundations. *IEEE Intelligent Systems*, 32(4):78–86, 2017.
70. Markus Rohde, Peter Brödner, Gunnar Stevens, Matthias Betz, and Volker Wulf. Grounded design—a praxeological research perspective. *Journal of Information Technology*, 32:163–179, 2017.
71. Tuomas Nissinen. User experience prototyping: a literature review. 2015.
72. Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, et al. Technology probes: inspiring design for and with families. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 17–24, 2003.



73. Jakob Nielsen and Rolf Molich. Heuristic evaluation of user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 249–256, 1990.