

Identifying predictors for energy poverty in Europe using machine learning

Willem van Hove^{a,b}, Francesco Dalla Longa^b, Bob van der Zwaan^{b,c,d,*}

^a Technical University of Munich, Department of Informatics, Munich, Germany

^b TNO, Energy Transition Studies, Amsterdam, the Netherlands

^c University of Amsterdam, Faculty of Science (HIMS and IAS), Amsterdam, the Netherlands

^d Johns Hopkins University, School of Advanced International Studies (SAIS), Bologna, Italy



ARTICLE INFO

Article history:

Received 25 January 2022

Revised 21 March 2022

Accepted 26 March 2022

Available online 31 March 2022

Keywords:

Energy poverty

Pan-European drivers

Household energy consumption

Explainable AI

Gradient boosting

ABSTRACT

In this paper we identify drivers for energy poverty in Europe using machine learning. The establishment of predictors for energy poverty valid across countries is a call made by many experts, since it could provide a basis to effectively target energy-poor households with adequate policy measures. We apply a “low income, high expenditure” framework that classifies households as being at risk of energy poverty to a dataset from a survey conducted at the household-level in 11 European countries with vastly different economies, cultures, and climates. A gradient boosting classifier is successfully trained on a set of socio-economic features hypothesized as predictors for energy poverty in this diverse set of countries. The classifier's internal model is analyzed, providing novel insights into the intricacies that underlie energy poverty. We find that besides the main driver - income - floor area and household size can be confirmed as predictors. Our results suggest the presence of universal predictors that are valid across Europe, and contextual ones that are governed by local characteristics. To facilitate advanced research into energy poverty in Europe, we recommend to increase and streamline household data collection efforts, both at the country- and EU-level.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The concept of energy poverty has been studied since the oil crises in the 1970s. In the United Kingdom, the term fuel poverty was coined; a person affected by fuel poverty was defined as “a person [who] is a member of a household living on a lower income in a home which cannot be kept warm at reasonable cost” [41]. Although fuel poverty is a field of active research, it is still poorly understood [42], and identifying fuel-poor households and targeting them with adequate policies remains a challenge [38,36,6]. Energy poverty is considered a similar, but broader concept than fuel poverty [4]. While fuel poverty is mainly associated with the affordability of (fossil) fuels to provide heating services in dwellings, the term energy poverty also encompasses the issues of energy efficiency, house insulation, electricity provision and fossil phase-out in the residential sector. In this paper we use the term energy poverty to refer to adequate provision of energy services at large in dwellings, in the context of the energy transition. Besides in the UK and Ireland, energy poverty research has largely

focused on developing countries, where access to modern energy services is not assured for a substantial share of the population. However, the transition to a more sustainable energy supply is expected to radically transform energy infrastructure, both in developing and developed countries. The cost of this transformation needs to be distributed in such a way that all households continue to be able to afford their energy bills. This has led to energy poverty research gaining more prominence in developed countries that are on the forefront of the energy transition as well. In these countries, energy poverty is defined as the inability of a household to afford its energy bills [10,4,6]. Issues of energy poverty in Europe may become particularly apparent in times of economic crisis or international conflict, such as currently experienced as a result of the war between Russia and Ukraine.

The field of energy poverty has become truly widespread in the European policy discourse since 2008, as in that year EU institutions and consultative committees began calling for a Europe-wide definition of energy poverty [40]. In 2013, the European Economic and Social Committee called for “European energy poverty indicators to be established and for statistics to be harmonised in order to identify, prevent and tackle the problem more effectively at the European level and to generate solidarity in this area” [9].

* Corresponding author.

E-mail address: bob.vanderzwaan@tno.nl (B. van der Zwaan).

Many indicators and metrics have been, and are being, proposed in the literature to measure energy poverty at the European scale (see e.g. [3], for a recent overview) and at the national and local scales (see e.g. [21,18]). A large-scale study in 2019 found that an indicator for energy poverty valid across countries is still absent. This has resulted in a lack of research comparing countries, since no metric exists that can be used to contrast them [20]. Concurrently, Castaño-Rosa et al. [6] observed a lack of standards to assess energy poverty across Europe and, instead, argued for a multiple indicator approach as starting point for policy decisions to reduce energy poverty in Europe. A recent study assumed that constructing all-encompassing predictors to assess energy poverty is unfeasible [36]. These three papers highlight (i) the need for the establishment of consistent energy poverty metrics that can enable cross-country comparisons and assist European legislation, and (ii) the current debate on whether the concept of universal energy poverty predictors is a meaningful one.

This research aims to improve our understanding of energy poverty in Europe by employing a data-driven approach based on machine learning (ML). We make use of two unique advantages offered by ML. First, ML models are able to find complex relations in large datasets that would require excessive amounts of manual labor using traditional statistical tests or running and evaluating standard regressions [1]. Second, by applying eXplainable Artificial Intelligence (XAI) methods one can gain significant insights on the intricate links between the inputs and outputs of ML models, which facilitates the understanding of sophisticated systems [27]. This was recently demonstrated, for example, when XAI methods enabled researchers to identify crucial predictive biomarkers of disease mortality briefly after the outbreak of the COVID-19 pandemic [43].

A similar methodology has already successfully been used to study energy poverty predictors in the Netherlands [10]. The present research aims at extending this approach to investigate energy poverty predictors within a larger geographic area, involving eleven countries in Europe. Finding predictors valid across Europe could aid the assessment of the prevalence of energy poverty in European countries, and subsequently assist in adequate policy design. We attempt to identify pan-European predictors through the use of XAI. In Section 2 the ML technique of gradient boosting is described, as well as the energy poverty framework used to classify households into energy poverty risk categories. In Section 3 the results of applying the classification framework to the dataset, and the resulting gradient boosting model are presented and analyzed. In Section 4 the findings are interpreted and the presence of two types of predictors, universal and contextual, is hypothesized. Furthermore, recommendations regarding the collection and accessibility of data are given, in order to better address energy poverty in Europe, and provide guidance for future research endeavors in this field.

2. Methodology

For this research, a survey on energy use in Europe conducted in 2018 by Enable-EU is used [12]. Enable-EU is an ongoing endeavor funded by the European Union's Horizon 2020 research and innovation program with the mission statement: "[Enable-EU] seeks to understand what determines people's choices in three key consumption areas: transportation, heating & cooling, and electricity" [13]. The survey was targeted at a group of 11 diverse countries in Europe: Bulgaria (BG), France (FR), Germany (DE), Hungary (HU), Italy (IT), Norway (NO), Poland (PL), Serbia (RS), Spain (ES), Ukraine (UA), and the United Kingdom (UK). A report outlining and comparing the outcomes of participating countries was published along with the dataset [19]. While some questions, for example

on prosumers, are country-specific, all respondents were asked to complete those sections of the survey with generic and socioeconomic questions, which provides us with a complete dataset on these topics. Most notably, participants were asked to report their income and energy expenditure, two variables crucial for assigning data points to the energy poverty risk classes employed in this paper. The diversity in the assessed countries and the homogeneity of the socioeconomic questions make this dataset particularly interesting for the purpose of investigating energy poverty predictors at the European level.

A more detailed characterization of the dataset, along with an in-depth explanation of the procedure we employed for data preprocessing can be found in van Hove [23]. Additionally, the energy expenditure and income distribution per country in the dataset is plotted in Figs. A.1 and A.2.

The energy poverty classification framework proposed by Dalla Longa et al. [10] is used to categorize each household in our dataset into one of four energy poverty risk categories. The framework operates on an income vs. energy expenditure grid that is divided into four quadrants using two thresholds, one for each axis. This is illustrated in Fig. 1. A household with income above and energy expenditure below the respective thresholds is categorized as "No risk" (green). If one of the variables crosses a threshold, the label "Income risk" (yellow) or "Expenditure risk" (orange) is assigned. If both thresholds are crossed, the household is categorized as being at highest risk of energy poverty, labelled "Double risk" (red).

The Enable-EU survey required respondents to classify their income after taxes and other deductions into the corresponding decile in their respective country. The brackets were presented as income ranges of the respective country the respondent lived in. The results can be categorized into 10 brackets. This enables a normalized measure of income across the participating countries. The income threshold is set for all countries between income deciles three and four. This particular threshold choice is made for three main reasons. First, analysis showed that for most countries the respective minimum wage corresponds to bracket 3. Second, the Low Income High Costs (LIHC) indicator applied in the UK, results in the vast majority of households categorized as energy poor being in the first three income deciles [35]. Third, this threshold ensures that all risk classes contain enough data points to produce reliable classifiers [10]. Energy expenditure is derived using two questions in which respondents were asked to report their latest heating and electricity costs. During preprocessing, all currencies are converted to Euros and all costs to annual costs, yielding yearly energy expenditure values for all households. The energy expenditure threshold is set at the 80th quantile of the absolute energy expenditure in the respective country, resulting in a different absolute value for every country. The 80th quantile threshold was chosen based on two main considerations: (i) this value was used in Dalla Longa et al. [10] to reproduce expected energy poverty levels found in prior studies for the Netherlands, and (ii) this choice ensures that the number of datapoints in the minority category is sufficient to train a ML model. While these two considerations provide a robust indication that the 80th quantile is adequate for our purposes, we recognize that this particular choice is somewhat arbitrary. As explored in detail in Dalla Longa et al. [10], this is linked to the practical difficulty of defining an objective threshold for energy poverty.

Mapping all households of the Enable dataset into the framework defined in Fig. 1 allows us to assign one of the four energy poverty risk labels to each datapoint. This labelled dataset allows us to employ supervised ML to train a model that attempts to predict the correct label for each individual household, represented by a selection of socio-economic features. For this purpose we employ an ML technique known as gradient boosting. This technique incre-




Fig. 1. Visual representation of the quadrants that the classification framework uses to assign households to a risk category.

mentally adds weak prediction models - in our case decision trees - to obtain a so-called ensemble that displays a better prediction performance than any of the weak models it is built from [17]. Gradient boosting achieves state-of-the-art performance on many modeling tasks [33], i.e. recommender systems. Especially on tabular data, decision tree methods consistently achieve high performance scores and provide a vast range of tools to analyze the internal model [27,8], which we can exploit to gain insight into the complex mechanisms linking energy poverty risk with its drivers. The Enable dataset contains many categorical features, i.e. one answer is to be selected from a list of possible options (categories). CatBoost is chosen as gradient boosting library since, as opposed to other popular gradient boosting libraries, it can deal with categorical features without requiring extensive preprocessing by the user [33,11].

The final step in our analysis entails assessing the identified features in terms of the influence they have on the final model outcome, i.e. estimating their feature importance. We use XAI techniques to quantify feature importance scores according to different metrics. These can be thought of as indicators of a feature's predictive power in the sense that, if a feature displays high importance scores for a good-performing model, this would suggest that it is a true predictor of energy poverty. A detailed description of the ML and XAI methodologies used in this work can be found in the supplementary online document and in van Hove (2020) [23].

3. Results

We divide the exposition of our results into three parts: (i) classification of the Enable dataset into our framework, (ii) the final ML model that was trained on the labeled data, (iii) analysis of the ML model with XAI techniques to identify and comprehend energy poverty predictors.

3.1. Energy poverty classification

The distribution over the income deciles in the entire dataset can be found in Fig. 2, left panel. Ideally, a representative sample of a population has the same number of respondents in each income bracket. However, in the Enable dataset a selection bias towards the lower-income brackets can be observed. On the country-level, this discrepancy is more prominent. This skewness

of the income distribution on the country-level is not deemed a fundamental problem for our purposes, as we are seeking predictors that are valid across Europe, and over the entire dataset the selection bias is less pronounced. The right panel of Fig. 2 presents the distribution of yearly energy expenditure from the Enable dataset. This varies with a wide spectrum from nearly zero to €¹5000 per household, with an overall median of € 997. At country-level Ukraine has the lowest per household median expenditure of € 267, and Norway has the highest median of € 2311. The majority of households in the dataset spend up to € 2000 per year on energy, with those spending more than that predominantly being from Norway and, to a lesser extent, from other Western European countries.

In order to provide a qualitative assessment of the chosen classification framework, in Fig. 3 we show the answers to a self-reported poverty question for the entire dataset. Each household is a colored line: the color depicts the response to the question "which of the descriptions below [sic] comes closest to how you feel about your household's income nowadays?". All households are plotted on a grid with income bracket on the x-axis and relative energy expenditure on the y-axis. The relative energy expenditure is an approximated estimate. Absolute energy expenditure is available in the data; however, income is only available in deciles. An approximate income for each household is obtained by taking the disposable income per decile for each country from the Statistics on Income and Living Conditions dataset [16]. This is used to provide a uniform measure of energy expenditure, thus allowing all households to be plotted on the same axis. Since households struggling to make ends meet are in general more vulnerable to energy poverty, we expect the distribution of colors in Fig. 3 to roughly match the four quadrants defined in Fig. 1. Indeed a gradient from red ("finding it very difficult on current household") in the top left, to green ("living comfortably on current income") in the bottom right can be observed in Fig. 3. To further highlight this, in Fig. 4 we apply our energy poverty classification framework to Ukrainian households. The qualitative resemblance of Fig. 3 and Fig. 4 corresponds to the close connection between poverty and energy poverty, and provides an empirical indication that the classification framework employed in this work can be effective in identifying energy poverty.

The energy poverty class composition that results from applying the risk framework on the dataset, for each of the 11 countries in the Enable-EU dataset, is plotted in Fig. 5. The countries are sorted by the share of households classified in the double risk category. The proportion of households classified as having the highest risk of energy poverty, i.e. the double risk category, is readily apparent from this figure (red portion of the bars). As the expenditure threshold is set at the 80th quantile, a horizontal line denoting this threshold between income risk and expenditure risk can be observed. We notice a distinction between Western European countries, characterized by lower double risk shares, and Central and Eastern European (CEE) countries, where double risk shares are more pronounced. The UK is found to be an outlier in this respect: while widely considered a Western European country, it has a significantly larger share of households classified in the double risk category than other Western European countries. We speculate that this apparent anomaly could partially be explained by the fact that the UK is the birthplace of capitalism, and to this day is considered closer to the free-state market than mainland Western European countries, generally regarded to have larger welfare states. Energy poverty is researched extensively in the UK, and several government policies aimed at reducing it have been implemented. This could be because disproportionately more

¹ All euros reported in this research are in €(2018).




Fig. 2. Income distribution of the whole dataset (left), and distribution of energy expenditure per household per year in the whole dataset (right).




Fig. 3. Poverty as reported by all households in the dataset.

households experience energy poverty in the UK than in other Western European countries. Energy poverty is shown to affect public health [25], and its prevalence in the UK might be linked with the high long-term excess winter mortality rates in this country [26].

In a perfectly representative sample of a country, the income risk and double risk group combined should add up to 30% as the income threshold is set between the third and the fourth income decile of a country. However, this is not always the case in our dataset, where the sum of the two groups ranges between 15% for Germany, to 82% for Poland. As we already observed, this is because the distributions of sampled households are skewed to the higher and lower income deciles, for Germany and Poland respectively. This is not deemed problematic to our main research goal, i.e. the identification of universal energy poverty predictors. It does, however, limit the conclusions that can be drawn regarding the prevalence of energy poverty in each of the countries, as this can also be caused by selection bias. The observed double risk shares in Fig. 5 are a consequence of two phenomena at play. On the one hand, we have genuine energy poverty prominence in a country, resulting in a large number of points in the double risk category. On the other hand, we have the non-representative sampling of income deciles in countries, resulting in misplaced thresholds and misrepresented risk groups. To what extent each factor plays a role, differs per country and falls beyond the scope of this

research. Consequently, although the data presented in this study provide a first indication for the prevalence of energy poverty in these countries, further research with improved data quality and quantity is necessary to confirm these observations. For comparison, data from previous research conducted by Dalla Longa et al. [10] in the Netherlands is also reported in Fig. 5 in the right-most bar of the plot [7]. In that study a similar energy poverty classification framework was applied. The relatively low and high shares of Dutch households, falling respectively into the double risk and no risk categories, lead us to conclude that the energy poverty profile of the Netherlands is close to that of other Western European countries. The extent to which this comparison is meaningful is, however, limited since the Dutch dataset is very different from the Enable dataset used in this research: it contains more data points, continuous numeric answers, and concerns neighborhood averages instead of single households.

3.2. Machine learning models

A CatBoost model is trained to classify households into one of the four energy poverty risk categories. A household is represented by seven selected features that can be found in Table 1. These features are hypothesized as potential predictors and selected as a result of extensive data analysis and domain expertise. Some features, such as floor area, are categorical but have a distinct ordering




Fig. 4. Energy poverty classification for Ukraine.




Fig. 5. Energy poverty classification distribution for countries in the Enable-EU dataset (the columns for the Netherlands is derived from Dalla Longa et al. [10]).

Table 1

Features used in the model, type, and number of unique values.

name	type	# unique values
income bracket	integer	10
floor area	integer	7
household size	integer	17
house detachment	integer	5
house age	integer	9
birth year respondent	integer	79
heating strategy	categorical	5

to them, and are represented by an integer. These features are therefore depicted as type “integer” and do not require any treatment before being used in a decision tree model. House detach-

ment corresponds to the survey question “Which best describes your home?” and has four possible answers. The answers range from “single-family house detached from any other house” to “apartment in a building with 6 or more flats,” and are interpreted as a scale that determines how well insulated the home is by surrounding homes. Multiple country-level studies have hypothesized that the level of house detachment plays a role in energy poverty [32,2]. Heating strategy corresponds to a question on what heating methods households employ. The feature value can be one of 5 different heating strategies described. This is a categorical feature with no clear ordering, and thus cannot be used directly in the gradient boosting method. It has type “categorical” and the necessary preprocessing steps are handled by the CatBoost library. Even

though income is both used to label and to classify the data, it is included as a feature to the ML model. Income is easy to obtain for policy makers and can thus be considered in policy design. Moreover, this will ensure that the found predictors are not surrogates of income, but instead of energy poor households.

Initially, the data is split in two datasets, a training and a test set. The test set is used exclusively to evaluate a model trained using the training set, on data the model has not seen before. In accordance to general Machine Learning best practices, the test set consists of 20% of the sample where each set has the same label distribution [1]. Subsequently, the training set is itself split into a training and a validation set in order to validate the performance of our model during training time. Here an 80/20 split was used as well. Additional details on the training, validation and test sets are provided in van Hove [23]. In Fig. 6 a confusion matrix is used to visualize the performance of our model, for the validation set (left panel) and the test set (right panel). A confusion matrix tabulates the labels as predicted by the model versus the true labels. In this figure, we have normalized all rows to sum to 100%. Values on the diagonal of a confusion matrix are known as the true positive rates: a model able to correctly identify all instances would result in 100% true positive rates for all categories. A model randomly assigning labels, i.e. random guessing, would result in a score of 25% in all cells of the confusion matrix. Given our classification framework (Fig. 1), using income as the only feature would allow a model to learn the income threshold and split the task into two binary classification tasks, i.e. discerning between double risk and income risk for datapoints below the income threshold, and distinguishing expenditure risk from no risk for datapoints above the income threshold. After this split, the performance would deteriorate to almost random guessing (50% diagonal).

The performance of such a one-feature model can be improved on by including the other features introduced in Table 1. Our resulting model performs much better than a 50% diagonal in the confusion matrix, yielding true positive rates between 60% and 74% on the test set. Similar scores were achieved on the validation set, indicating that the model is not just memorizing the training data - known as overfitting - but that the performance generalizes well to unseen data. The true positive rates vary per category. Compared to a similar study of the Netherlands, where diagonals scores ranging between 73% and 82% were achieved [10], the results achieved in this study are worse.

3.3. Explainable Artificial Intelligence analysis

Having successfully built an ML model to classify energy poverty risk in a diverse group of European countries, we can now assess to what extent the various features included in the model influence its outcome. In other words, we attempt to open the black box that is the ML model and gain insight into its inner workings [34]. In order to accomplish this we introduce two measures of feature importance: The permutation importance and the mean absolute SHapley Additive exPlanation (SHAP) value. The values of these two metrics for all seven features in our model are plotted in Fig. 7. For both methods, the feature importance values are normalized, such that the highest value equals 100 and the rest of the values are given relative to that one. As income bracket is clearly the most important feature for both methods, the x-axis is broken to allow easy comparison of the importance score of the other features. Permutation importance assigns an importance score to a feature based on the effect of shuffling its values on the model performance [5]. For each feature, the data are permuted several times, each time yielding a (slightly) different score. The resulting deviation between runs is depicted by the error bars in the plot. SHAP values have a strong mathematical foundation in cooperative game theory [28]. For a single datapoint, each feature's contribution to the final prediction of the model is computed as if it were a coalition game in which each feature would get a "payout". This is done for each individual datapoint in the unseen test set. The mean absolute value of all contributions is used to attribute each feature an importance score. No error bars are obtained with this metric.

Both feature importance measures identify income bracket, floor area, and household size as the main drivers of our model. The features house age and birth year respondent are assigned the two lowest importance values in both metrics. For permutation importance, the error bars span from negative to positive values. This implies that these are not good predictors for the model and thus can be disregarded as such. This is confirmed by the findings of the mean absolute SHAP values, that also indicate they are of little importance, although not as overtly as with permutation importance. With respect to heating strategy and house detachment, the results are inconclusive at this stage. Mean absolute SHAP value assigns heating strategy a slightly higher importance than house detachment. However, permutation importance results




Fig. 6. Confusion matrix of the final model, on the validation (left) and test (right) set.




Fig. 7. Permutation and mean absolute SHAP value feature importance of the seven features in our model plotted with a broken x-axis.




Fig. 8. Partial dependence plots for the seven features of our model.

rank heating strategy as having no impact on the model performance with error bars running from negative to positive values: further research is required to resolve this divergence.

A way to assess how a feature value affects the model output is by using Partial Dependence Plots (PDPs). Friedmann proposed PDPs as a comprehensive summary of the model dependence in his paper introducing gradient boosting [17]. The PDPs for each of the risk classes and each of the features are plotted in Fig. 8. It shows the partial dependence between the model's output for every energy poverty risk group and every single feature of the model, averaging over all other features. The y-axis depicts the partial probability of the category values on the x-axis. A straight line suggests that the feature has no influence on the model's prediction. A very volatile line, i.e. a line for which y values change significantly when x values move from one category to the next, suggests that the feature has a big impact on the model outcome. This is the case, for example, for the lines in the income bracket panel (top left) in Fig. 8. The shared y-axis enables a simple comparison of the volatilities of the partial dependencies. PDPs, while providing some valuable insights, do not give a complete explanation of how the features affect the model's output because interactions between features are averaged out to get this 2-dimensional visualization.

The PDP for income confirms what can also be observed in the confusion matrices. Based on income bracket, the model correctly reduces the multiclass classification task to two binary ones: if the income is below the income threshold, probabilities for the no risk and expenditure risk categories drop to zero. If the income is above the threshold the inverse is true, and the probabilities of income risk and double risk drop to zero. In the other plots we can distinguish two pairs of classes with similar profiles. For each risk class pair, one exceeds the expenditure thresholds and can be considered the higher risk class of the two. The plots indicate that a larger house floor area increases the likelihood of a household being classified in the higher risk category corresponding to its income. The PDP for the feature household size shows a similar profile, where a larger household shows increased probabilities of being at risk. The PDPs of house age and birth year respondent are almost flat lines, confirming our results that did not regard them as predictors. House detachment and heating strategy - for which the feature importance results are inconclusive - show some volatility, again indicating that these might have some predictive power, but require further research to be established as predictors.

In order to provide additional insight in the internal workings of our ML model, we introduce the notion of a decision plot. A decision plot [29] exploits the additive nature of the feature contribution assigned by SHAP values by providing an effective visual summary. A visualization for a single household from the test set can be seen in Fig. 9. The y-axis has the features followed by the feature value in brackets. The effect of the feature is depicted between the two adjacent horizontal lines. The values on the x-axis correspond to the model output, these values are transformed by the model to probabilities, which are depicted in brackets in the legend. The dashed line is the eventual classification done by the model. This plot depicts how a large household living in a large home, detached from other houses, living of an income in the fifth decile, is correctly classified as being in the expenditure risk group. The split of the task into two binary classification tasks can clearly be observed (red and yellow versus green and orange lines). The little impact the features heating strategy, birth year respondent, and house age have on the prediction is also apparent. By analyzing these plots, one could in principle differentiate several characteristic decision paths corresponding to certain types of households that are more prone to be in energy poverty. We believe this could prove to be beneficial for policy design in the future. In order to properly carry out this type of analysis, however, one would need a larger and more consistent dataset than the one our current model is based upon.

4. Conclusion and discussion

We have successfully applied ML to find features that have demonstrated their predictive power in a heterogeneous dataset comprised of 11 countries, representative of the European continent. The resulting model is better at classifying households categorized as being at risk of energy poverty than a theoretical model with only income as an input. However, it does not attain the same level of accuracy as a gradient boosting model achieved in a similar study that addressed one single country, i.e. the Netherlands [10]. This leads us to hypothesize the presence of two types of predictors for energy poverty in Europe. Universal predictors are drivers valid across a varied set of countries, e.g. for the whole of Europe. Besides these, there are also predictors concerned with the local specificities of a country (or region), which we call contextual predictors. Among the drivers considered in this study, examples of




Fig. 9. Decision plot of a single sample in the test dataset. The feature values are in brackets.

universal predictors are income – low income households have a higher chance of incurring energy poverty – and floor area – energy consumption is generally higher in larger dwellings. An example of a contextual predictor is house age, given that the correlation between building quality and age is country-dependent. Universal predictors can serve as a starting point for European countries to establish an overarching framework to assess energy poverty. These can then be complemented in individual countries (or regions) by contextual predictors, thus adequately assessing the prominence of energy poverty at (sub)national level. This concept of having a national definition of energy poverty that is complementary to a common European definition has been proposed by other studies in the literature [40], our research supplies empirical support to this notion.

Our energy poverty classification framework is essentially a Low Income High Costs (LIHC) scheme. An analogous framework is currently used in the UK for energy poverty monitoring. Many other quantitative and qualitative frameworks have been proposed in the literature and applied in practice. Some well-known examples are the 10% quote – which classifies as energy poor all households whose energy expenditure is above 10% of their income – and the 2 M scheme – which selects households that spend more than twice the national median on their energy bills. All existing frameworks have their particular advantages and drawbacks, and the quest for adequate energy poverty indicators and metrics remains a very active field of research. For our study we chose a quantitative framework that is currently employed in practice, and that is suited to the available data and that lends itself to be analyzed with a ML classifier.

The features income, floor area, and household size are all found to be of significant importance to our energy poverty risk classifier. Our results suggest these three features as universal predictors on the European continent. While our results are inconclusive with regards to the features house detachment and heating strategy, the features house age and birth year of respondent have been found to be insignificant on a supranational scale. House age has previously been found to be an indicator of energy poverty [22]. We assign this apparent discrepancy to the heterogeneous nature of the dataset used. As was established by Galev et al. [19], house age cannot be considered a proxy for insulation on a supranational level. Nonetheless, it likely acts as a contextual predictor. One can in principle hypothesize – and find evidence for – the existence of other energy poverty drivers, in addition to those considered in our study. These include, for example, energy prices, dwelling energy efficiency, ownership status, number of unemployed persons in households. Our choice is based on (i) drivers that were readily available in the present dataset, (ii) drivers that are expected to be easily found or measured for future application of our method, and (iii) drivers that were already identified as possible energy poverty predictors in previous studies.

The results reported in this paper should be considered in the light of certain limitations. Absolute energy expenditure is used to classify households into an energy poverty risk group, as the dataset does not allow for easy conversion to relative energy expenditure. As a result, some households might be misclassified as energy poor. Furthermore, households that severely under-consume (a phenomenon also known as hidden energy poverty) are missed by this framework; this atypical form of energy poverty is an open research question and beyond the scope of this research. Future research could use an improved EP classification framework and attempt to recreate the findings of this paper. Furthermore, this research limits its focus to socio-economic indicators of energy poverty. An interesting research direction would be one where spatial features, such as nightlights or surface temperature, are consid-

ered as this has been found to affect energy consumption [30]. The arbitrariness of choosing the 80th quantile of the absolute energy expenditure as a threshold is a reflection of the practical problem of finding adequate boundaries to define energy poverty for policy purposes. Alternative threshold choices would lead to selecting different slices of the population, i.e. different socio-economic groups. Only by studying these in detail and validating the results with interviews and surveys can one hope to gain a more objective view on what an adequate set of energy poverty thresholds may be.

The limited size of the Enable database might have caused our model to underfit, resulting in a classifier too simple to capture the full complexity of the modeling task. As noted, the Enable dataset also suffers from selection bias with respect to the income deciles. While this was not deemed to significantly affect our results, more homogeneity across income distributions would be desirable in order to achieve EU-wide models that are also representative at national level, i.e. to capture the full predictive power of universal energy poverty drivers. The lack of high-quality data regarding energy poverty has long been recognized in the field [31,37,39,2,20,24], and several initiatives have been launched to attempt to make progress in this respect, e.g. the Energy Poverty Observatory (EPOV) was founded with (among others) the goal to improve and harmonize data collection [14]. Based on our findings, we endorse this type of initiatives and we recommend that additional efforts to collect and make available consistent high-quality datasets both at national and at EU level are pursued. Future research might apply the approach introduced in this paper to such high-quality datasets to effectively train and analyze more sophisticated ML models. We argue that ML methods can be used as an efficient means to examine and evaluate different energy-poverty metrics, and highlight non-obvious and complex correlations between different drivers.

To stimulate the energy transition to a more sustainable energy supply, some sort of carbon taxation is imposed in many states in the EU and US. Moreover, as part of the European Green Deal, the EU is launching the Just Transition Fund to enable regions to address the impacts of the energy transition [15]. Spending the aforementioned funds on energy R&D and technology innovation is a commonly heard suggestion. Another suggestion is to use those revenues to assist people in meeting energy poverty challenges, which may be exacerbated due to stringent climate change mitigation measures. One question raised is how to determine which people will deserve and receive assistance, and which do not. The approach proposed in this paper, or an improved version thereof, can help in determining who may be in dire straits when it comes to affording basic energy services, and ultimately help to identify appropriate types of assistance for different groups of energy poor households. Provided the availability of large representative high-quality datasets, XAI methods could be instrumental in efforts to distribute greenhouse gas taxes revenues through policy, if the targets include the alleviation of (energy) poverty and establishment of equity among consumers of energy services.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix

Figs. A1 and A2 present the distributions of, respectively, energy expenditure and income from the Enable-EU dataset.




Fig. A1. Absolute (top), and approximate relative energy expenditure (bottom) for all countries in the data set.




Fig. A2. Income distributions of respondents per country.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.enbuild.2022.112064>.

References

- [1] E. Alpaydin, *Introduction to Machine Learning*, The MIT Press, 2014.
- [2] C.A. Bollino, F. Botti, Energy poverty in Europe: A multidimensional approach, *PSL Quarterly Review* 70 (2017) 473–507, https://doi.org/10.13133/2037-3643_70.283_4.
- [3] S. Bouzarovski, H. Thomson, M. Cornelis, Confronting energy poverty in Europe: A research and policy agenda, *Energies* 14 (4) (2021), <https://doi.org/10.3390/en14040858>.
- [4] S. Bouzarovski, Energy poverty in the European Union: Landscapes of vulnerability, *Wiley Interdiscip. Rev. Energy Environ.* 3 (2014) 276–289, <https://doi.org/10.1002/wene.89>.
- [5] L. Breiman, Random forests, *Machine Learn.* 45 (2001) 5–32, <https://doi.org/10.1023/a:1010933404324>.
- [6] R. Castaño-Rosa, J. Solís-Guzmán, C. Rubio-Bellido, M. Marrero, Towards a multiple-indicator approach to energy poverty in the European Union: A review, *Energy Build.* 193 (2019) 36–48.
- [7] CBS Statistics Netherlands, 2018, StatLine publicaties Kerncijfers wijken en buurten. <https://www.cbs.nl/nl-nl/dossier/nederland-regiona>. Accessed 21-07-2020.
- [8] Chen, T., Guestrin, C., 2016. XGBoost: A scalable tree boosting system, in: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, New York, NY, USA. pp. 785–794. URL: <http://doi.acm.org/10.1145/2939672>. 2939785, doi:10.1145/2939672.2939785.
- [9] P.J. Coulon, H. Bataller, Opinion of the European Economic and Social Committee 'for coordinated European measures to prevent and combat energy poverty', *Off. J. Eur. Union C* 341 (2013) 21–27.
- [10] F. Dalla Longa, B. Sweerts, B. van der Zwaan, Exploring the complex origins of energy poverty in The Netherlands with machine learning, *Energy Policy* 156 (2021), <https://doi.org/10.1016/j.enpol.2021.112373>.
- [11] A.V. Dorogush, V. Ershov, A. Gulin, CatBoost: gradient boosting with categorical features support, in: *NIPS 2017*, 2018, pp. 1–7, URL: <http://arxiv.org/abs/1810.11363>, arXiv:1810.11363.
- [12] Enable-EU, 2018. D4.1 dataset for the comparative sociological analysis of the household survey results. http://www.enable-eu.com/wp-content/uploads/2019/10/enable_eu_dataset_households.zip. Accessed 01-04-2020.
- [13] Enable-EU, 2020. The context. <http://www.enable-eu.com/about-us/>. Accessed 24-06-2020.
- [14] European Commission, 2020a. Energy Poverty Observatory. www.enenergypoverty.eu/.
- [15] European Commission, 2020b. Proposal for a regulation of the European parliament and of the council establishing the just transition fund. COM 22 final.
- [16] Eurostat, 2018. Distribution of income by quantiles - eu-silc and echp surveys (ilc_di01). Data retrieved from Eurostat 06-2020, https://ec.europa.eu/eurostat/web/products-datasets/-/ILC_DI01.
- [17] Friedman, J.H., 1999. Greedy Function Approximation: a Gradient Boosting Machine. *Annals of Statistics* 29, 1189–1232. URL: <http://www.jstor.org/stable/10.2307/2699986>.
- [18] I. Faiella, L. Lavecchia, Energy poverty. How can you fight it, if you can't measure it?, *Energy Build* 233 (2021) 110692.
- [19] Galev, T., Csd, A.G., Pellerin-carlin, T., 2018. D4.1 Final report on comparative sociological analysis of the household survey results. Technical Report. Enable-EU.
- [20] Hassani, H., Yeganegi, M.R., Beneki, C., Unger, S., Moradghaffari, M., 2019. Big Data and Energy Poverty Alleviation. *Big Data and Cognitive Computing* 3, 50. doi:10.3390/bdcc3040050.
- [21] J.P. Gouveia, P. Palma, S.G. Simoes, Energy poverty vulnerability index: A multidimensional tool to identify hotspots for local action, *Energy Rep.* 5 (2019) 187–201, <https://doi.org/10.1016/j.egyr.2018.12.004>.
- [22] J.D. Healy, J.P. Clinch, Quantifying the severity of fuel poverty, its relationship with poor housing and reasons for non-investment in energy-saving measures in Ireland, *Energy Policy* 32 (2) (2004) 207–220.
- [23] W. van Hove, Identifying Energy Poverty Predictor Using Machine Learning, *Technical University Munich*, 2020, Master's thesis.
- [24] L. Karpinska, S. Śmiech, Invisible energy poverty? Analysing housing costs in central and eastern Europe. *Energy Research & Social, Science* 70 (2020), <https://doi.org/10.1016/j.erss.2020.101670>.
- [25] C. Liddell, C. Morris, Fuel poverty and human health: A review of recent evidence, *Energy Policy* 38 (2010) 2987–2997, <https://doi.org/10.1016/j.enpol.2010.01.037>.
- [26] C. Liddell, C. Morris, H. Thomson, C. Guiney, Excess winter deaths in 30 European countries 1980–2013: a critical review of methods, *J. Public Health* 38 (2015) 806–814, <https://doi.org/10.1093/pubmed/fdv184>.
- [27] S.M. Lundberg, G. Erion, H. Chen, A. DeGrave, J.M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, S.-I. Lee, From local explanations to global understanding with explainable ai for trees, *Nat. Machine Intell.* 2 (1) (2020) 56–67.

- [28] Lundberg, S.M., Lee, S.I., 2017. A unified approach to interpreting model predictions, in: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (Eds.), *Advances in Neural Information Processing Systems* 30. Curran Associates, Inc., pp. 4765–4774. URL: <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>.
- [29] S.M. Lundberg, B. Nair, M.S. Vavilala, M. Horibe, M.J. Eisses, T. Adams, D.E. Liston, D.K.W. Low, S.F. Newman, J. Kim, S.I. Lee, Explainable machine-learning predictions for the prevention of hypoxaemia during surgery. *Nature, Biomed. Eng.* 2 (2018) 749–760, <https://doi.org/10.1038/s41551-018-0304-0>.
- [30] B. Mashhoodi, D. Stead, A. van Timmeren, Land surface temperature and households' energy consumption: Who is affected and where?, *Appl Geogr.* 114 (2020) 102125.
- [31] S. Pachauri, D. Spreng, Measuring and monitoring energy poverty, *Energy Policy* 39 (2011) 7497–7504, <https://doi.org/10.1016/j.enpol.2011.07.008>.
- [32] L. Papada, D. Kaliampakos, Measuring energy poverty in Greece, *Energy Policy* 94 (2016) 157–165.
- [33] Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A.V., Gulin, A., 2018. Catboost: Unbiased boosting with categorical features. *Advances in Neural Information Processing Systems*, 6638–6648 arXiv:1706.09516.
- [34] Ribeiro, M.T., Singh, S., Guestrin, C., 2016. "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*, Association for Computing Machinery, New York, NY, USA. p. 1135–1144. URL: <https://doi.org/10.1145/2939672.2939778>, doi:10.1145/2939672.2939778.
- [35] J.C. Romero, P. Linares, X. López, The policy implications of energy poverty indicators, *Energy Policy* 115 (2018) 98–108, <https://doi.org/10.1016/j.enpol.2017.12.054>.
- [36] S. Sareen, H. Thomson, S. Tirado Herrero, J.P. Gouveia, I. Lippert, A. Lis, European energy poverty metrics: Scales, prospects and limits, *Global Transitions* 2 (2020) 26–36, <https://doi.org/10.1016/j.glt.2020.01.003>.
- [37] S. Scarpellini, P. Rivera-Torres, I. Suárez-Perales, A. Aranda-Usón, Analysis of energy poverty intensity from the perspective of the regional administration: Empirical evidence from households in southern Europe, *Energy Policy* 86 (2015) 729–738.
- [38] B.K. Sovacool, Fuel poverty, affordability, and energy justice in England: Policy insights from the Warm Front Program, *Energy* 93 (2015) 361–371, <https://doi.org/10.1016/j.energy.2015.09.016>.
- [39] H. Thomson, S. Bouzarovski, C. Snell, Rethinking the measurement of energy poverty in Europe: A critical analysis of indicators and data, *Indoor Built Environ.* 26 (2017) 879–901, <https://doi.org/10.1177/1420326X17699260>.
- [40] H. Thomson, C.J. Snell, C. Liddell, Fuel poverty in the European Union: a concept in need of definition?, *People, Place Policy online* 5–24 URL (2016) <http://eprints.whiterose.ac.uk/105609/>.
- [41] UK Parliament, 2000. Warm homes and energy conservation act. <https://www.legislation.gov.uk/ukpga/2000/31/contents>.
- [42] R. Walker, P. McKenzie, C. Liddell, C. Morris, Area-based targeting of fuel poverty in Northern Ireland: An evidenced-based approach, *Appl. Geogr.* 34 (2012) 639–649, <https://doi.org/10.1016/j.apgeog.2012.04.002>.
- [43] L. Yan, H.T. Zhang, J. Goncalves, Y. Xiao, M. Wang, Y. Guo, C. Sun, X. Tang, L. Jing, M. Zhang, X. Huang, Y. Xiao, H. Cao, Y. Chen, T. Ren, F. Wang, Y. Xiao, S. Huang, X. Tan, N. Huang, B. Jiao, C. Cheng, Y. Zhang, A. Luo, L. Mombaerts, J. Jin, Z. Cao, S. Li, H. Xu, Y. Yuan, An interpretable mortality prediction model for COVID-19 patients, *Nature Machine Intell.* 2 (2020) 283–288, <https://doi.org/10.1038/s42256-020-0180-7>.