ESRI Working Paper No. 762

September 2023

Energy poverty prediction and effective targeting for just transitions with machine learning

Constantine Spandagos[1,2*], Miguel Angel Tovar Reaños[1,2] & Muireann Á. Lynch[1,2]

a) Economic and Social Research Institute, Dublin, Ireland
b) Department of Economics, Trinity College Dublin, Dublin, Ireland

*Corresponding Author:
Dr Constantine Spandagos
Economic and Social Research Institute,
Whitaker Square, Sir John Rogerson's Quay,
Dublin, Ireland
Email: constantine.spandagos@esri.ie

Keywords: energy poverty prediction; energy poverty targeting; machine learning; EU member states; just energy transitions

JEL codes: D10, I30, Q40, Q41, Q48

Energy poverty prediction and effective targeting for just transitions with machine learning

ABSTRACT

The prevalence of energy poverty as a major challenge in numerous countries, the escalating energy crisis that generates supply fears and increased prices, and the need to build just supporting mechanisms within the net zero energy transition add impetus to improving our ability to accurately predict energy vulnerable households. In Europe, this is hindered by limited recognition of the fact that energy vulnerable households are not necessarily income poor (and vice versa). Artificial Intelligence, and machine learning techniques in particular, may be applied to improve the efficient targeting of energy poverty schemes, enabling the accurate prediction of energy vulnerable households via objective, publicly available data. However, such applications are still limited, especially across a large number of countries. In response to the above, we develop an innovative machine learning framework for accurate prediction and fair targeting of energy poor households across all the current members of the European Union and the United Kingdom. While we explore various machine learning algorithms, most of our analysis is performed using a Random Forest classifier. Our approach to explore energy poverty beyond income reveals household-level and country-level predictors of energy poverty, such as dwelling condition, energy efficiency, social protection payments and gas supplier switching rates. We also demonstrate how machine learning algorithms can offer straightforward visualization of the mechanism that determines the energy poor classification, enabling alleviation schemes to be transparent and assisting policy-makers in setting more effective thresholds for assistance allocation. Furthermore, we evaluate the potential fairness of alleviation schemes and demonstrate that basing their targeting solely on income-relevant or social welfare-relevant criteria would be ineffective and result in significant numbers of energy poor households being excluded from energy assistance.

Keywords: energy poverty prediction; energy poverty targeting; machine learning; EU member states; just energy transitions

JEL codes:
D10
I30

Q40
Q41
Q48

1. Introduction

Worldwide, numerous governments have pledged to meet the climate targets of the Paris Agreement by making their jurisdictions carbon-neutral (Rogelj et al., 2016). For instance, China, the United States (US) and the European Union (EU), the world's three largest carbon dioxide (CO_2) emitters (Ortega-Ruiz et al., 2022), have set targets to create net zero emission economies within the next few decades. As the energy sector is a major contributor of worldwide CO_2 emissions (Marcucci and Fragkos, 2015), decarbonizing the manner in which energy is generated, distributed and used is at the forefront of carbon neutrality efforts. However, certain energy decarbonization policies, if not coupled with appropriate supporting mechanisms, may have regressive effects, for instance by limiting the ability of vulnerable populations to access and afford energy services- especially as the current global energy crisis (von Homeyer et al., 2022) creates fears for fuel supply and further increases in energy costs (Steckel et al., 2022). These developments are realized against a background of existing income inequalities (Johns et al., 2013) combined with procedural and distributional injustices pertaining to energy availability, access and affordability (Carley and Konisky, 2020). Consequently, the energy transition may provide real social benefits only if it progresses in an equitable and just manner, accompanied with measures that effectively alleviate the phenomenon of energy poverty (González-Eguino, 2015), which is broadly defined as the inability of households to meet their energy needs.

Mitigating energy poverty is a challenging task because the phenomenon is multifaceted, with various possible causes, dimensions and manifestations. At the same time, there is no widely-accepted framework to define it and understand it (Roberts et al., 2015). The traditional approach was to explore energy poverty solely in relation to income, but recent research is demonstrating that energy poverty has multiple and hidden dimensions (Cong et al., 2022). In the EU, while member states are urged to mitigate energy poverty, the current directives and suggested guidelines do not specifically address households that are energy vulnerable; instead, they predominantly focus on households at risk of poverty in the financial sense only (Kyprianou et al., 2019). This is a serious shortcoming in the EU policy framework, as protecting income poor consumers and addressing energy poverty are distinct challenges that require distinct solutions (Pye et al., 2015b; Tovar, 2021). In fact, it has been suggested that one of the greatest energy poverty alleviation challenges will be to properly identify those who deserve to receive assistance (Dubois, 2012). Recognizing the distinction between energy poor and income poor consumers and constructing an EU-wide knowledge base on what constitutes energy vulnerability in households would be crucial steps in strengthening the European energy transition with effective and fair alleviation strategies. Assuming the goal of any energy poverty alleviation strategies is to identify and assist households that cannot afford to meet their energy requirements, a framework that enables policy-makers to identify households that are unable to meet their energy requirements via objective, publicly available data is clearly advantageous. Such a framework would facilitate the efficient targeting of alleviation strategies without relying on self-reported household energy deprivation.

To contribute to that cause, we develop a modeling framework for the prediction and targeting of energy poverty across all the countries of the EU-28 group (i.e., the current 27 member-states and the United Kingdom), while departing from the traditional approach of understanding the phenomenon in terms of income only. We determine a mapping to energy vulnerability not only from household income, but also from various other household-level and country-level predictors, with a particular focus on building features, energy efficiency, the countries' social protection policies and their energy market characteristics. In so doing, we provide a data-driven way of mapping from objective household and country-level data to a subjective report of energy poverty. To predict energy poverty based on this diversity of potential predictors and to understand their relative importance, we first build a unique dataset merging information from four reputable sources of European statistics covering the years 2010-2020. We include variables that are normally not included previous analyses, such as switching fuel supplier rates. Subsequently, we use approximately half a million data points from our dataset to train various machine learning algorithms. As a subset of Artificial Intelligence (AI) applications, machine learning models have been proven capable of handling complex tasks (Department of Business Energy & Industrial Strategy (BEIS), 2017; Hong and Park, 2021; Levi, 2021; van Hove et al., 2022), and increasingly popular in energy-relevant research. However, their application to the field of energy poverty alleviation, and in particular to evaluating the targeting effectiveness of policies, is still limited (López-Vargas et al., 2022).

While we explore several machine learning algorithms, most of our analysis is performed using a Random Forest model due to its ability to provide the most accurate predictions across data previously unseen by the system. To the best of our knowledge, this is the first analysis that combines state of the art machine learning techniques and the best available data to identify household and country-specific drivers of energy poverty, and quantify the consequences of current policies to tackle this condition across the whole EU-28 group. Our analysis reveals that apart from income, the condition of the dwelling inside which a household lives, a country's social protection payments and gas supplier switching rates are the most important predictors of energy poverty within the examined variables. By identifying the configuration that generates accurate predictions based on testing data for 2 time periods (2010-2019 and 2020), we demonstrate its usefulness is dealing with new batches of data, previously unseen by the system, without the need to rely on self-reported energy poverty data, and without the need of being re-trained. Finally, and to the best of our knowledge, this work is the first to demonstrate how machine learning algorithms can be the basis of designing energy poverty alleviation schemes that contribute to procedural and distributional energy justice by being transparent and fair. Specifically, we suggest the use of Decision Tree algorithms when the main objective is to offer transparency on who is receiving energy assistance and Random Forest algorithms when the objective is to evaluate the schemes' fairness potential more accurately. For the latter purpose, we hypothesize energy poverty alleviation schemes based exclusively on income-based metrics or social welfare parameters and conceptualize their fairness potential as their effectiveness in ensuring that no households in need of energy assistance are excluded from it due to improperly set criteria. Our analysis demonstrates that 16-56% of households that would have been non-recipients based on income thresholds, typical energy poverty metrics, householders' age or unemployment status are in need of energy assistance. This observation reveals the weakness of frameworks solely based on such criteria in properly addressing energy poverty across Europe. At the same time, it demonstrates the potential of our approach in guiding the development of more effective, EU-wide energy poverty supporting policies.

The remainder of this paper is organized as follows. Section *2* provides a brief overview of energy poverty and the application of AI techniques on understanding it. Section *3* details the paper's methods and describes the mathematical principles behind Random Forest, the main machine learning algorithm used herein. Subsequently, Section *4* presents the paper's results, in terms of evaluating the selected model performance, the prediction variables' importance and the fairness of hypothetical alleviation schemes. Section 5 discusses the findings and their policy implications, while providing closing remarks.

2. Background

This section provides a brief review of the latest energy poverty literature. During the last decade, the majority of scholarly and policy work on this issue was concerning the United Kingdom (UK) and Ireland. However, there is today an abundance of energy poverty research conducted in numerous countries. This makes energy poverty a vast topic; therefore, this literature review is not meant to be exhaustive but rather to provide a brief overview of the energy poverty subtopics that are most relevant to this paper, namely its definitions, measurements, drivers, the application of machine learning methods to understand it, and challenges associated with its targeting. Standard methods (Spandagos et al., 2021) were employed to gather a sample of representative scientific publications on these subtopics, such as: searching well-established electronic databases using keywords including (but not limited to) "energy poverty", "fuel poverty", "artificial intelligence", "machine learning" and "targeting", and; refining the results through multiple title/abstract reviews.

2.1 Energy poverty: terminology, drivers and metrics

Even though they had been defined differently in the past, it is now common for the terms "energy poverty" and "fuel poverty" to be used interchangeably. According to traditional understanding, "energy poverty" concerns the lack of access to modern energy services in developing countries, while "fuel poverty" refers to the affordability of such services in developed countries (Bouzarovski and Petrova, 2015). It has been suggested that more work on defining the scope of the two concepts is required to properly and mutually integrate them (Li et al., 2014). In practice, however, "energy poverty" is gradually becoming a prevailing common term for describing the situation of households not being able to meet their energy needs in developed countries as well, such as the US (Bednar and Reames, 2020), Japan (Okushima, 2016) and member states of the EU (Kyprianou et al., 2019). For instance, the majority of official EU policy documents and recent directives employ the term "energy poverty" (Pye et al., 2015b). For the remainder of this paper, "energy poverty" will also be used as the umbrella term to broadly define the inability of households to meet their energy needs due to limited supply, affordability, quality, quantity, reliability, or a combination of the above (Cong et al., 2022).

Towards theorizing global energy poverty through a vulnerability lens, Bouzarovski and Petrova (2015) have developed a framework that categorizes the drivers of energy vulnerability into six groups, namely to drivers that are associated with i) access, ii) affordability, iii) flexibility, iv) energy efficiency, v) needs and vi) practices. The authors of that work recognize the limited availability of energy carriers to appropriately meet household needs as a key access-relevant factor. In terms of affordability, a major driver for energy vulnerability is considered to be the high ratio between fuels and household income- with tax or assistance schemes included. Another driving factor in that category is the inability to invest in new energy infrastructure. The inability to switch to a more favorable energy services provision is a driving factor that falls in the flexibility-relevant category, and the high loss of useful energy during domestic conversion is an efficiency-relevant driver. As for the needs-relevant category, energy vulnerability is associated with the imbalance between a household's energy needs and the available services. Finally, a major driver in the practice-relevant category is the lack of knowledge of ways to mitigate energy poverty, such as participating in assistance-providing programs and using energy more efficiently. Given the multiple dimensions of drivers, it is no surprise that energy poverty has multidimensional consequences that extend, among others, to public health, gender roles and educational opportunities (Sovacool, 2012).

There is no universally accepted standard for determining whether an individual (or a household) is energy poor (Roberts et al., 2015). Instead, a variety of measures, with their own strengths and weaknesses, have been employed. Cong et al. (2022) categorizes energy poverty metrics into primary or secondary and relative or absolute. Primary metrics directly employ consumer-level information, while secondary metrics use weighted scoring of certain indices and aggregated information from utilities. Furthermore, absolute metrics measure energy poverty via strict thresholds, while relative metrics provide comparative information across multiple households, countries, or regions. Another common categorization is to distinguish energy poverty metrics as subjective and objective. The most common example of primary-relative and subjective metrics comes directly from households and individuals and concerns answers to survey questions about their ability to meet their energy needs. On the other hand, the most common example of an objective and primary-absolute metric concerns the percentage of income that should not be exceeded for paying for energy services. Pioneering work in the UK has set the maximum acceptable threshold for this purpose to be 10% of the available income (the so-called "10% rule") (Dogan et al., 2022). Another objective metric is the Low Income High Cost

(LIHC) indicator (Robinson et al., 2018), which classifies households as energy poor when their energy expenditure is higher than the median in their country, and at the same time their income after that expenditure falls below the country's poverty line (Siksnelyte-Butkiene et al., 2021). The 10% rule and the LIHC indicator are commonly used in developed countries, as energy poverty is widely measured there using economic (González-Eguino, 2015) and income thresholds. However, while income is a crucial component of energy poverty, it is not adequate to explain all the dimensions of energy poverty. Energy poor households are not necessarily income poor households (and vice versa) and income-based measurement approaches do not adequately incorporate factors such as equipment efficiency, type of household, or type of fuel used. Therefore, one of the most common criticisms against the 10% rule is that it is often unable to distinguish between the energy poor and the income poor. For instance, employing that metric might lead to the misclassification as energy poor of households that have both high income and high energy consumption (Hills, 2012). Similarly, the LIHC indicator has been criticized because it might misclassify as non-energy poor low-income households that simply consume less energy than a given threshold. Apart from economic thresholds, physical thresholds have also been used to measure energy poverty in relation to the minimum energy consumption that can satisfy basic needs (González-Eguino, 2015). The shortcoming in that approach is that determining what constitutes a "basic need" can be subjective and not independent of local culture or lifestyle. Moreover, energy poverty has been measured in relation to technological thresholds when the focus is on accessing modern types of energy (such as electricity) instead of sources like firewood, coal, kerosene and dung (Wang et al., 2021). Finally, given its diverse nature, energy poverty is frequently measured on the basis of multidimensional indicators, such as the ones suggested by Pachauri et al. (2004) and Nussbaumer et al. (Nussbaumer et al., 2013).

2.2 Machine learning applications

AI is considered one of the most innovative breakthroughs in modern technology (López-Vargas et al., 2022). Machine learning is a subfield of AI that enables computers to learn and perform processes without being explicitly programed to do so (Samuel, 1959). Machine learning algorithms can be constructed and trained through sample data, and then be employed to make predictions. They have demonstrated the ability to make accurate predictions, handle atypical relations between variables (Hong and Park, 2021) and identify complex relations in large datasets, in settings where it would normally require excessive amount of manual labor and time if traditional statistical methods were used (van Hove et al., 2022). Because of these advantages, and despite their limitation of often operating as "black boxes" (when there is limited visibility of how the decision rules of the algorithms are formed), machine learning models are currently very popular in handling complex tasks such as image recognition, fraud detection and medical diagnosis (BEIS, 2017).

The number of examples of applying AI techniques, including machine learning techniques, for the specific purpose of guiding energy poverty alleviation is still limited (López-Vargas et al., 2022); however, the trend for conducting work in that direction is growing. The majority of existing work focuses predominantly on identifying the most influential energy poverty predictors in one or more countries. A recent example of work focusing on developed countries is the study by Dalla Longa et al. (2021) on identifying energy poverty predictors in the Netherlands. The authors of that work employed machine learning to classify Dutch households into four energy poverty risk categories, and identified variables such as income, household size, and dwelling age, value and ownership as important predictors. van Hove et al. (2022) extend the geographic scope of that work to explore energy poverty predictors in 11 European countries. The latter work demonstrated that income, household size and floor area are more universal predictors of energy poverty, while the effects of dwelling age depend more on local conditions. In another European example, the UK's Department of Business Energy & Industrial Strategy (BEIS, 2017) employed machine learning to predict the energy poverty status of British households. Hong and Park (2021) explored several machine learning algorithms to predict energy poverty in South Korea, and identified income, floor area, food expenses, and householder age and education as important predictors.

Pertaining to examples of studies focusing on developing countries, Wang et al. (2021) found that a combination of machine learning with geographical and environmental remote-sensing data predicts energy poverty more accurately than approaches based on socio-economic indicators only. More recently, Abbas et al. (2022) employed an multidimensional energy poverty index (MEPI) approach to explore extreme energy poverty in 59 developing countries in Asia and Africa. The authors of that work identified which countries within that group were the most vulnerable to extreme multidimensional energy poverty, as well which socio-economic factors are the strongest predictors of energy poverty. Apart from a household's accumulated wealth, the marital status of the householder, the size and ownership of the residence and the location of the residence emerged as other important factors.

Apart from the small but growing number of machine learning studies focusing specifically on energy poverty, there is a plethora of machine learning applications focusing on other aspects that are, however, closely related to energy poverty and may assist efforts to better understand it. López-Vargas et al. (López-Vargas et al., 2022) provide a detailed review of research in that direction, with a particular focus on aspects such as income, building energy consumption and performance, energy prices, energy billing irregularities and thermal comfort. In terms of exploring income-relevant variables, the authors of that review provide examples of applications such as: predicting income levels in the US (Chakrabarty and Biswas, 2018), studying urban gentrification in the UK (Reades et al., 2019), predicting the risk of suffering social exclusion in Spain (Serrano et al., 2018), predicting socio-economic indicators in Germany (Feldmeyer et al., 2020), monitoring housing rental prices in China (Hu et al., 2019) and detecting unemployment in Ireland (Curbelo Montañez and Hurst, 2020). Pertaining to the exploration of building energy performance, examples are the applications of machine learning to improve feature engineering of energy data mining (Zhang et al., 2018) and to identify the most influential features affecting energy use intensity (Ma and Cheng, 2016), both using US data. Other applications include predicting electricity prices in Germany (Uniejewski et al., 2019), identifying non-technical losses and irregular energy consumption based on US and Indian data (Sharma et al., 2017) and predicting thermal comfort perceptions of the elderly using data collected in China (Wang et al., 2019).

Most of the studies mentioned in this subsection were conducted within the last 5 years, indicating that applying machine learning techniques in the fields relevant to energy poverty is a timely approach, which attracts growing interest. Furthermore, the wide range of national settings used as case studies in the aforementioned work denotes the international character of that interest. Machine learning techniques that were employed in the aforementioned studies include, but are not limited to Gradient Boosting, k-Nearest Neighbors and Random Forest. The latter is a technique that has become particularly popular in topics related to energy and climate policy and is described in detail in Section 3 of this paper.

2.3 Challenges for EU policy and the current paper

The EU has developed various policies to alleviate energy poverty, either in the form of directives or through funding for country-level initiatives. A detailed overview of the latest relevant directives is provided by Kyprianou et al. (2019). Examples of existing country-level interventions include Italy's "social bonus" discount mechanism and the UK's Warm Homes Discount program, while France has relied (or plans to rely) on social tariffs and energy vouchers that may also be used to cover efficiency upgrades. However, regardless of the type of assistance offered (financial assistance, energy efficiency improvements, information provision, or other), previous programs have been criticized not only because of their dubious targeting, but also because of improperly identified eligibility thresholds. One of the possible explanations behind the inability for proper targeting is the fact that the limited recognition of the energy poor-income poor distinction has been projected into EU legislation. For instance, the directives on common rules for the electricity and gas markets consider "energy poor" to be equal to "vulnerable consumer". This creates additional difficulties in establishing a widely accepted framework for alleviating energy poverty, as member states have their own national criteria to determine what constitutes a "vulnerable" household. For instance, the UK, Spain and Portugal define vulnerable households based on criteria including old age and eligibility to receive benefits from the social welfare system (Boardman, 2009; Sareen et al., 2020). Such criteria may have little connection with one's actual status of being energy poor or vulnerable, as indicated by their lack of association with energy poverty indicators suggested by the Energy Poverty Observatory (EPOV) for statistical purposes (Sareen et al., 2020).

The literature focusing on the targeting of energy assistance programs is still limited (Best et al., 2021). Dubois (2012) discusses different practical approaches for identifying energy poor households in France, including the employment of direct identification via database crossing and geographical identification. It should be noted that the latter work positions targeting as the first step in a three-step process for energy poverty policy, with the other two steps being identification and implementation. Consequently, targeting is defined to concern the choice of the population that should benefit from a policy, and the political and economic feasibility of that choice. Subsequently, identification concerns selecting the process and criteria based on which energy poor households will be identified. In practice, however, the literature often employs the term "targeting" to encompass both steps. The concept of geographical identification is also examined by the works of Walker et al. (2012) and of Gupta and Gregg (2018), where Geographic Information Systems techniques were used to explore energy policy targeting in the UK. Earlier work by Raffio et al. (2007) explored the targeting of energy efficiency retrofits using weather and bill data collected from student residences in the US. Reames (2016) focused on community-based targeting for energy efficiency schemes in the same country. More recently, Best et al. (2021) and Wang and Lin (2022) applied multidimensional indicator approaches to examine energy poverty targeting in Australia and China, respectively. The authors of the former work argued for the usefulness of employing wealth-based instead of income-based criteria in the targeting process, while the latter work found a polarization between income poverty and energy poverty that excludes a large number of households from receiving assistance.

While work in these directions broadens our knowledge about critical targeting concepts, the generally limited size of the energy-relevant targeting literature and the small number of studies that apply AI to targeting challenges might be hindering a deeper understanding of how energy poverty alleviation schemes can be improved. In an example of improper targeting outside the energy poverty field, Andini et al. (2018) demonstrated that €7 billion in monetary gains could have been achieved if machine learning methods were used to inform tax rebate programs and properly identify the consumption-constrained households in Italy. It is not yet fully known what monetary amounts would have been saved or properly allocated if energy poverty alleviation schemes in European countries were based on machine learning tools too. We found the task of identifying applications of machine learning techniques with the purpose of evaluating energy poverty targeting to be challenging. There is previous machine learning work on the evaluation of policies with an energy essence, such as the one by Yin and Zhou (2022) on evaluating China's photovoltaic poverty alleviation policy. However, work such as the latter focuses mostly on evaluating a policy's ability to deliver social, environmental and economic benefits, rather than on targeting issues. We conclude that machine learning applications on the effective targeting of energy poverty schemes is unexplored in the literature to date, and that there is also a dearth of relevant models that are built by combining data from a wide range of countries.

Improper targeting at the country-level contributes to improper targeting at the EU-level too. In turn, the lack of an EU-level understanding of how to effectively identify energy poor consumers hinders the development of an authoritative body and structured EU-wide strategy that would be fully dedicated to energy poverty alleviation (Kyprianou et al., 2019). An important step in that direction is the development of the Energy Poverty Observatory (EPOV) as a joint space to facilitate new policy development (Kyprianou et al., 2019), but the need remains for the development of an expanded, EU-wide knowledge base about the most crucial components of effective energy poverty prediction and targeting.

2.4 Contribution to the literature

In view of the above, and to contribute to the construction of such knowledge base, we develop a machine learning framework for the prediction and targeting of energy poverty across all countries of the EU-28 group. We base this on a unique and large database consisting of merged household-level and country-level European data from various sources. To the best of our knowledge, this work is the first to develop machine learning classification mechanisms that can jointly: i) predict energy poverty across all the EU-28 countries; ii) provide transparent visualizations of the underlying assigning mechanisms that determine whether an individual is classified as energy poor or not- a functionality that overcomes the "black box" limitation of similar applications and can contribute to the pursuit of procedural energy justice; iii) evaluate the fairness of energy poverty schemes whose targeting relies on income-based or social welfare criteria, by quantifying their effectiveness in ensuring that none of the energy poor households is excluded from assistance- a functionality that can contribute to the pursuit of distributional energy justice- and; iv) integrate both household-level and country-level predictors of energy poverty in the targeting evaluation process. In the light of lacking a standardized framework to compare energy poverty across countries, the latter integration of predictors from two levels is useful in building a pan-European understanding of energy poverty targeting. In short, this paper provides a mapping from a subjective measure of energy poverty, self-reported ability to heat one's home, to an objective, cross-country dataset, in a data-driven manner. This approach is not withstanding the existing literature which shows that subjective and objective evaluations of energy poverty, while closely correlated, may in fact diverge.

While some of the objectives of our paper (particularly (i) and (ii) above) are addressed via MEPIs, we propose our machine learning approach as a complementary approach, for several reasons. The major advantage of MEPIs is that they correct for bias that arises from computing energy poverty as a function of income only, as does our approach. The advantages of our approach over MEPI include the fact that MEPIs require a (subjective) weighting of various metrics, including but not limited to ability to keep a dwelling sufficiently warm, the fact that the data to apply the machine-learning approach are readily available from the EU Statistics on Income and Living Conditions (EU-SILC) database and can be applied to future waves of SILC data, and the fact that 64.4% of energy use in EU households is for space heating (Eurostat, 2023a). Thus, our machine learning approach computes one metric that reflects the majority of energy end use in the EU. The approach also identifies, in a data-driven way, the household and country characteristics that contribute to energy poverty (as measured by ability to adequately heat one's home), while the identification and weighting of those characteristics are often chosen by the researcher under MEPIs.

For the remainder of this paper, targeting is defined to encompass the first two steps of the process described by Dubois (2012). It is assumed that the group selected to benefit from energy assistance policies is all the energy poor households of a given jurisdiction, and that it is politically feasible to do so. Therefore, the first step defined by Dubois (2012) is complete and the use of the term targeting herein predominantly concerns the step of identifying the energy poor households. Consequently, the terms "targeting households" and "identifying households" are often used interchangeably in the sections that follow. The scope of our analysis does not include the implementation step, but we focus on the improvement of the targeting/identification step, as we consider it to be a critical step towards effective implementation.

3 Data and Methods

3.1 Data

To build this study's data base, we merged household-level and country-level data from four reputable sources of EU-wide statistics. The scope covered 10 years of data from all the EU-28 countries. In total, an initial sample approaching 2.5 million observations was built. Table 1 summarizes information about all the variables involved.

Our microdata were obtained from the EU-SILC surveys (Eurostat, 2023b) for all the EU-28 countries and for the years 2010-2020. EU-SILC surveys collect cross-section and longitudinal data on income, poverty, social exclusion and living conditions and are a primary source of such data for Eurostat, the EU's statistical office. Detailed descriptions of the data collection methodology behind EU-SILC surveys are provided by Eurostat. The survey results provided identification numbers for both individual and households, it was thus possible to identify individuals living in the same household. This enabled us to initially merge all individual data to clusters of household data. In each EU-SILC survey, a particular question was asking the respondents to indicate whether they have the ability to keep their home adequately warm or not. The answers to that question determined this study's dependent variable, namely the household's energy poverty status, and determining a mapping from objective data to this variable is the major aim of this research. Consequently, the output variable of the models that were developed in the remainder of this paper reflects perceived energy poverty. Figure 1 displays the percentages of households that identified themselves as energy poor in each EU-28 country according to the surveys.

In terms of selecting the predictors, i.e., the independent variables that potentially predict the energy poverty status, the intention was to depart from the traditional approach of understanding energy poverty in relation to income only. Therefore, apart from household income, other household-level variables were selected for investigation. The majority of these concerned building characteristics such as: the type of the dwelling, the number of rooms that are available to the household for living, and the condition of the dwelling in terms of damages and darkness. Pertaining to damages, a variable was included to assess whether the dwelling has leaking roof, damp walls/floor/foundations and rot in window frames or floors; pertaining to darkness, a variable was included whether the dwelling is too dark, without adequate natural light. The remaining household-level characteristics that were selected for examination included the type of household and possible cash benefits (monetary support) that the household has received. Data concerning all these variables were acquired from the EU-SILC surveys.

Finally, several country-level predictor variables were obtained from Eurostat, Odyssee-Mure and the Agency for the Cooperation of Energy Regulators (ACER). The variables were chosen to reflect a country's: i) social protection policies and ii) heating fuel market characteristics. Hence, the selected variables included a country's heating energy efficiency levels achieved for households, gas prices for households and gas supplier switching rates (measured as the percentage of people within a country who have made such switch). The final variable that was included reflected a country's monetary assistance provided to its citizens in the form of social protection payments.

Table 1: Summary of all variables involved in this study.

Variables	**Source**	**Type/ Coding**	**Mean**	**Std. dev.**
A. Dependent:				
Energy poverty	EU-SILC	Binary: '1' is energy poor (is not able to keep home adequately warm), '0' is non-energy poor (is able to keep home adequately warm)	0.11	0.31

B. Predictor:

B1) Household-level				
Household income (annual)*	EU-SILC	Continuous, in €	9.98	1.07
Household type	EU-SILC	Categorical: 9 categories following the EU-SILC categorization: ‘5’, one person household; ‘6’, 2 adults under 65 years, no children; ‘7’, 2 adults at least one older than 65, no children; ‘8’, other, no children; ‘9’, single parent, one or more children; ‘10’, 2 adults, 1 child; ‘11’, 2 adults, 2 children; ‘12’, 2 adults, 3 or more children; ‘13’, 2 other with children.	7.69	2.55
Dwelling type	EU-SILC	Categorical: 4 categories following the EU-SILC categorization: ‘1’, detached house; ‘2’, semi-detached house; ‘3’, apartment in building with less than 10 dwellings; ‘4’, apartment in building with more than 10 dwellings.	2.28	1.23
Damaged dwelling	EU-SILC	Binary: ‘1’ is damaged (leaking roof, damp walls/floor/foundation, rot in window frame/floor), ‘0’ is not damaged	0.15	0.36
Dark dwelling	EU-SILC	Binary: ‘1’ is dark (not enough natural light), ‘0’ is not dark	0.05	0.23
Number of rooms available to household	EU-SILC	Continuous	3.78	1.39
Cash benefits received by household	EU-SILC	Continuous, in €	3,105	17,509
B2) Country-level				
Social protection payments	Eurostat	Continuous, in € per capita	6,126	4,004
Gas prices	Eurostat	Continuous, in € per gigajoule	15.62	6.14
Increase in energy efficiency level for heating in households	Odyssee- Mure	%, change since the year 2000	31.45	22.30
Gas supplier switching rate	ACER	%	5.20	6.26
Observations: 2,419,500				

*We used logarithm of this variable in the regression analysis.

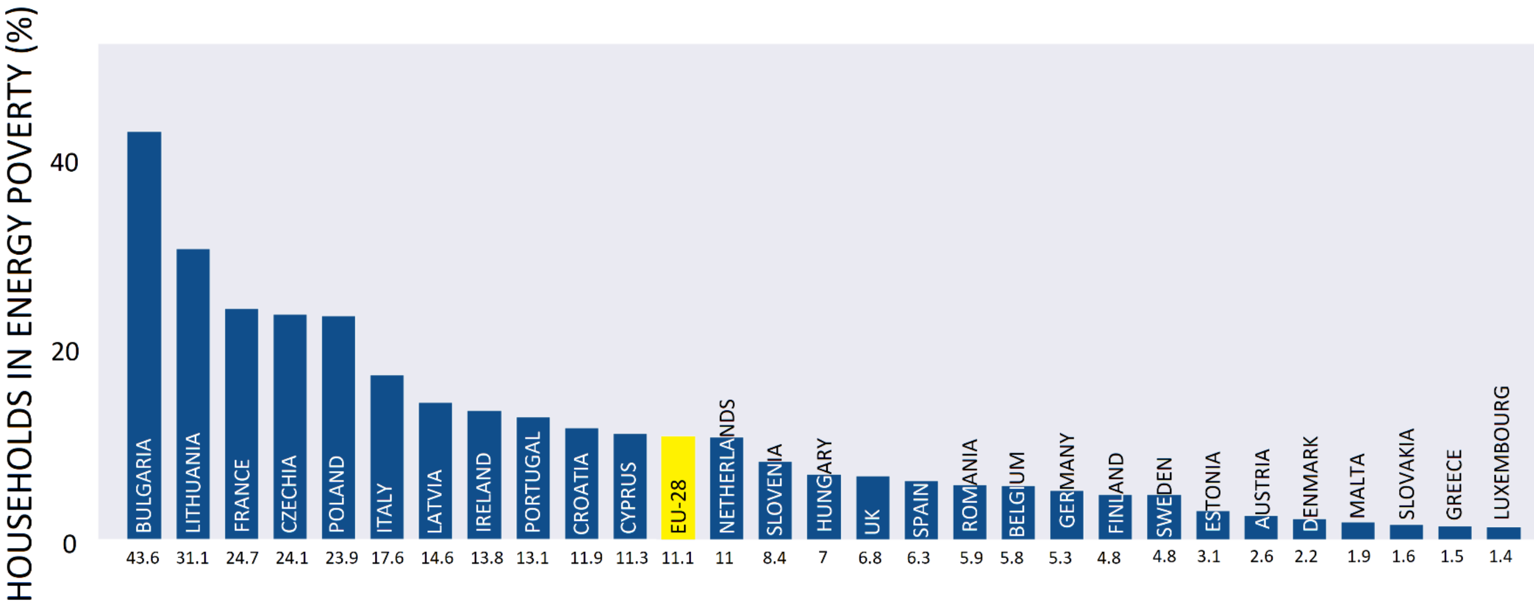

Figure 1: Percentages of households in (perceived) energy poverty per EU-28 country, according to responses to an EU-SILC survey question over the years 2010-2020. The question concerned the households' self-declared ability to keep their home adequately warm.

3.2 Modeling approach

We treated the problem of predicting the energy poor as a classification problem, where the target was to classify each observation (household) either as energy poor (the value of the output value is 1) or as non-energy poor (the value of the output variable is 0). As Figure 1 indicates, the average percentage of the (self-declared) energy poor households across the 28 European countries is 11.1%. Consequently, the original dataset was heavily dominated by observations of the non-energy poor class. Such disparity has been demonstrated to affect negatively the accuracy of classification models, as it distorts the decision boundary between classes (BEIS, 2017). To overcome this, and following standard practice for such cases (BEIS, 2017), the original dataset was down-sampled by removing a random selection of observations of the non-energy poor class. This resulted in a balanced (equal proportions of energy poor and non-energy poor observations) dataset of approximately half a million observations. According to standard machine learning practices, the next step should include splitting the data to training and testing datasets, with the purpose of the former being to train the model and of the latter to evaluate it based on data that are not part of the training process (van Hove et al., 2022). However, one prior step was executed here, i.e., to split the balanced set of half a million observations into 2010-2019 data and 2020 data. That was dictated by the need to keep the 2020 data intact, so that we can evaluate the targeting of alleviation schemes based on the latest available data. Furthermore, that splitting serves the purpose of exploring the model's ability to make predictions of comparable accuracy for different time periods (see *4.1*). After the split in 2020 and pre-2020 datasets, the latter was further split to training (80%) and testing (20%) subsets.

One of the most crucial steps in solving machine learning problems is to determine the algorithm that best meets the problem's objectives. Here, to achieve an initial understanding of which algorithm has the ability to generate accurate predictions based on this work's datasets, we employed the Tree-based Pipeline Optimization Tool (TPOT) (Le et al., 2020), a genetic programming tool used to determine machine learning algorithms that provide the highest accuracy given specific datasets (BEIS, 2017). Several runs of TPOT revealed Random Forest classifiers as the preferred machine learning algorithm. This was confirmed empirically at a later stage, by comparing the prediction accuracy of Random Forest classifiers of different configurations, with that of other popular machine learning algorithms (see Table 2 below). The following subsection provides an overview of Random Forest algorithms, which form the basis for the majority of this paper's analysis.

3.3 Random Forest algorithms

Random Forest is a powerful technique for supervised machine learning. In brief, a Random Forest model builds a collection of multiple decision tree algorithms (hence the term "forest") and merges them to generate an accurate and stable prediction, in regression and classification problems alike. The building of trees, or "estimators", is realized through the "bagging" mechanism (Strobl et al., 2009), an ensemble meta-algorithm that increases the accuracy of the final model by combining information from multiple learning models. Each estimator consists of three components: a single "root" node, the decision nodes and the "leaf" nodes. The nodes represent attributes that are used to predict the outcome. In the classification case, Random Forest classifiers follow a tree-shaped process: starting from the root node, they perform binary splits based on certain criteria until a leaf node is attained. The leaf node cannot be split any further and represents a final binary result.

Due to its ability to provide highly accurate predictions, Random Forest is becoming increasingly popular in fields including (but not limited to) genetics, clinical medicine, bioinformatics and psychological research (Strobl et al., 2009). Random Forest models are also gaining momentum in the fields of environment, climate and sustainability, in particular in exploring people's environmental attitudes (Beiser-McGrath and Huber, 2018), climate change awareness (Lee et al., 2015) and perceptions of climate-relevant policies (Levi, 2021). Pertaining to the machine learning studies on energy poverty (and closely-related aspects) mentioned in subsection *2.2*, a large number of them employed Random Forest models, either exclusively or together with other techniques (Bienvenido-Huertas et al., 2021; Curbelo Montañez and Hurst, 2020; BEIS, 2017; Feldmeyer et al., 2020; Hong and Park, 2021; Hu et al., 2019; Ma and Cheng, 2016; Reades et al., 2019; Serrano et al., 2018; Wang et al., 2019, 2021; Zhang et al., 2018).

Biau and Scornet (2016) provide a detailed description of the mathematical principles underlying Random Forest algorithms, both within regression and classification frameworks. In the former framework, a random input vector $\mathbf{X} \in [0,1]^p$ is assumed, and a random response $Y \in \mathbb{R}$ is predicted through a function $m(\mathbf{x})= \mathbb{E}\,[Y\,[\mathbf{X}=\mathbf{x}]]$ (Biau and Scornet, 2016). The target is to construct an estimate $m_n : [0,1]^p \rightarrow \mathbb{R}$ of the m function using a training dataset $D_n=(\mathbf{X}_1, \mathrm{Y}_1),\ldots, (\mathbf{X}_n, \mathrm{Y}_n)$ consisting of independent random variables that follow the same distribution as the pair $(\mathbf{X},Y)$. A random forest is an ensemble of M randomized estimators; for the j-th estimator in the ensemble, $m_n(\mathrm{x}; \Theta_j, D_n)$ denotes the predicted value at query point $\mathbf{x}$, with $\Theta_1,\ldots, \Theta_M$ being random variables, independent of D_n and distributed in the same manner as a random variable Θ. The purpose of the latter variable is to resample the training set before the development of individual estimators and to select the directions for splitting. The estimators are combined to form the forest following equation (1):

$$m_{M,n}(\mathbf{x}; \Theta_1, \ldots, \Theta_M, D_n) = \frac{1}{M}\sum_{j=1}^{M} m_n\big(\mathbf{x}; \Theta_j, D_n\big). \tag{1}$$

In the classification framework, m_n denotes a "classifier", a function of x and D_n that aims to predict the label Y, which takes values in $\{0, 1\}$. The classifier is considered to be consistent if its probability of error, denoted by $\boldsymbol{L(m_n)} = \mathbb{P}[\boldsymbol{m_n(X)} \neq \boldsymbol{Y}|\ \boldsymbol{D_n}]$, satisfies equation (2):

$$\lim_{n\to\infty} \mathbb{E}\, Lm_n = L^* \ , \tag{2}$$

where L^* is the error of the optimal Bayes classifier:

$$m^*(\mathbf{x}) = \begin{cases} 1, & \text{if } \mathbb{P}[Y = 1 \mid \mathbf{X} = \mathbf{x}] > \mathbb{P}[Y = 0 \mid \mathbf{X} = \mathbf{x}] \\ 0, & \text{otherwise} \end{cases} . \quad (3)$$

The Random Forest classifier is obtained through a majority vote among all the estimators:

$$m_{M,n}(\mathbf{x}; \Theta_1, \dots, \Theta_M, D_n) = \begin{cases} 1, & \text{if } \frac{1}{M}\sum_{j=1}^{M} m_n\left(\mathbf{x}, \Theta_j, D_n\right) > 1/2 \\ 0, & \text{otherwise} \end{cases} . \quad (4)$$

An important feature of Random Forest algorithms is their ability to rank the importance of variables of regression and classification problems through two measures of significance, namely the mean decrease impurity and the mean decrease accuracy (Biau and Scornet, 2016). The latter, also known as permutation importance (*PI*), is defined as the decrease in the accuracy of a Random Forest model when a single variable value is randomly permuted (Breiman, 2001). *PI* is often preferred, as it does not inflate the importance of numerical variables- something that is possible with mean increase impurity.

Assuming $\mathbf{X}=(\mathbf{X}^{(1)},\dots, \mathbf{X}^{(p)})$ and a forest of M estimators, the *PI* of variable $X^{(j)}$ is measured by randomly permuting the values of $X^{(j)}$ in the out-of-bag cases and by averaging the difference in the out-of-bag error estimation before and after the permutation over all estimators:

$$\widehat{PI}(X^{(j)}) = \frac{1}{M}\sum_{l=1}^{M}[R_n\left[m_n\left(\cdot, \Theta_l, D_{l,n}^{j}\right)\right] - R_n[m_n(\cdot, \Theta_l), D_{l,n}]] , \quad (5)$$

where $m_n(\cdot, \Theta_l)$ is the l-estimator estimate, $D_{l,n}$ the out-of-bag testing dataset of the l-th estimator, $D_{l,n}^{j}$ the same dataset with permuted values of $X^{(j)}$ and R_n is defined by:

$$R_n[m_n(\cdot, \Theta_l), D] = \frac{1}{|D|}\sum_{i:(\mathbf{X}_i, Y_i)\in D}^{M}(Y_i - m_n(\mathbf{X}_i, \Theta_l))^2 . \quad (6)$$

The population version of $\widehat{PI}\left(X^{(j)}\right)$ takes the following form:

$$PI^{*}(X^{(j)}) = \mathbb{E}\big[\mathrm{Y} - m_n(\mathbf{X}'_j, \Theta)\big]^2 - \mathbb{E}[Y - m_n(\mathbf{X}, \Theta)]^2 , \tag{7}$$

where $\mathbf{X}'_j = (X^{(1)}, X'^{(j)}, \dots, X^{(p)})$, with $X'^{(j)}$ denoting an independent copy of $X^{(j)}$. In the classification framework, equation (5) stands with $R_n(m_n(\cdot, \Theta_l), D)$ denoting the number of points that are correctly classified by $m_n(\cdot, \Theta_l)$ in dataset D (Biau and Scornet, 2016). The importance of all the predator variables involved in this study was estimated on the basis of the *PI* metric (see Figure 3).

4. Results and discussion

4.1 Model selection and performance

Confusion matrices are tables that summarize a classification model's overall accuracy, which is defined as the ratio of correct predictions to the total number of predictions made. The matrices also provide information about the True Positives Rate (TPR) and the False Positives Rate (FPR) (Jiao and Du, 2016) The former denotes the portion of the positive class (here, the energy poor class according to the self-reported measured gathered by the EU-SILC surveys) that was classified correctly, while the latter denotes the proportion of the positive class that was classified incorrectly. Corresponding measures for the negative class (here, the non-energy poor class) are the True Negatives Rate (TNR) and the False Negatives Rate (FNR), respectively. Note that FNR= 100- TPR and FPR= 100- TNR. We were primarily interested in the accuracy achieved for the energy poor class of the 2020 dataset, as our aim was to train a model that evaluates the targeting of energy poverty alleviation schemes based on the latest available data. At the same time, we wanted the model to make comparably accurate predictions using the pre-2020 dataset too. This is particularly important in adding confidence to the model being a readily available solution that performs well with future batches of data that are previously unseen by the system, without the need of re-training. Table 2 summarizes the performance of Random Forest models of different configurations, in terms of generating accurate predictions across the two datasets. The percentages displayed in Table 2 are taken from the corresponding confusion matrices of the models. Apart from accuracy, Table 2 indicates each model's F1 score (Jiao and Du, 2016) expressed as a percentage, as an additional measurement of overall performance. The F1 score is the harmonic mean of precision and recall; the latter (equal to TPR) is the number of true positives divided by the sum of true positives and false negatives, while the former is the number of true positives divided by the sum of true positives and false positives.

Two of the key parameters to be configured in Random Forest models is the number of estimators (trees) comprising the forest, and the maximum depth of each estimator. Table 2 indicates that increasing the maximum depth from 3 to 9 generally increases the models' accuracy and F1 scores. The same is true while increasing the number of estimators from 1 to 50. On the other hand, having a smaller number of estimators and a lower value of the maximum depth parameter provides the benefits of less complexity and faster computational execution of required functions. In view of that trade-off between accuracy/F1 score and complexity, a Random Forest specification with 10 estimators and the value of maximum depth set to 6 was empirically selected as optimal for the majority of this paper's analysis. It was empirically found that estimators with more than 6 layers were impractical to visualize and their assigning mechanisms difficult to interpret. Furthermore, and as Table 2 indicates, moving from 1 to 10 estimators exhibited a more obvious effect on improving accuracy and F1 score, while moving from 10 to 50 estimators did not improve them significantly. That Random Forest model generated a prediction for the study's class of central focus, i.e., the energy poor class of the 2020 data, with an accuracy approaching 80%. Furthermore, the model demonstrated an average (between both 2020 data classes) prediction accuracy of 72%, which is one of the highest prediction accuracies across all classes and datasets examined. Past work in similar settings has demonstrated the reliability of machine learning models with prediction accuracies in these ranges (for instance, Andini et al. (2018) and Wang et al. (2021)). The particular Random Forest model is indicated as model number 5 on Table 2 and herein referred to as our "selected" Random Forest model.

Subsequently, the performance of Random Forest models was evaluated against those of three other popular machine learning algorithms, namely Decision Tree, Extreme Gradient Boosting (XGBoost) (Nobre and Neves, 2019) and k-Nearest Neighbors (Kuncheva, 1995). Decision Tree is another algorithm for supervised machine learning that is based (as the name suggests) on the concept of decision trees and a similar structure with root, decision and tree nodes. It is, however, less complex than Random Forest as it constitutes of only one estimator instead of multiple. XGBoost is a form of Gradient Boosting (Friedman, 2001), another machine learning algorithm that employs decision trees, while k-Nearest Neighbors is a non-decision tree algorithm. Table 2 indicates that pertaining to this study' datasets, Random Forest generally outperforms k-Nearest Neighbors, which predicts the energy poor class of the 2020 data particularly poorly. On the other hand, the performance of XGBoost models is comparable to that of Random Forest models of equivalent specifications. Therefore, XGBoost can be considered a reliable alternative for future work with similar datasets. In this study however, the selected Random Forest model still generated more accurate predictions (by approximately 10%) for the energy poor class of the 2020 data, compared to an XGBoost model with the same number of estimators and maximum depth. The same can be said for Decision Tree models, which also exhibit comparable performance in terms of accuracy. However, Decision Tree is prone to overfitting, as it often works particularly well with training data but fails to accurately predict testing data. Random Forest avoids overfitting thanks to its reliance on multiple estimators (Breiman, 2001). This is important when the target is to generate accurate predictions on new and previously unseen data, as is the case in the current work. Moreover, among all the models examined, our selected Random Forest model exhibited the highest F1 score with the 2020 data and one of the highest with the 2010-2019 data. It should be noted that the highest F1 score with the 2010-2019 data was exhibited by k-Nearest Neighbors (with 2 "neighbors"). However, the large difference between that value and the corresponding one with the 2020 data does not create confidence that the particular model can produce results of comparable accuracy when different reference years are used. To further explore the ability of our selected Random Forest to produce predictions of comparable quality independently of the batch of data used, we run the model with other years' datasets too. Table A.1 of the Appendix indicates that the accuracies and F1 scores of the model do not exhibit noteworthy variations when other years' datasets are used.

Finally, Table 2 demonstrates the performance of a Random Forest model with similar specifications to the selected one (i.e., maximum depth set to 6, with 10 estimators), but consisting of household-level explanatory variables only. For both data sets, the accuracy and the F1 score of that model (indicated as model number 9 on Table 2) are smaller than those of the selected one. This indicates the importance of including country-level predictors in the modeling process.

Table 2: Performance comparison of Random Forest, Decision Tree, k-Nearest Neighbors and XGBoost models of different specifications and for 2 sets of testing data. All percentages are derived from the confusion matrices of the corresponding models. "Accuracy" is defined as the ratio of correct predictions to the total number of predictions made; correct predictions of the energy poor and non-energy poor classes are based on the TPR and TNR metrics,

respectively; F1 scores are also indicated. The “selected” Random Forest model, i.e., the one selected for the majority of this paper’s analysis is boldfaced and indicated as model number 5. The performance of that model if only household-level variables (and not country-level ones) were taken into account is also indicated (model number 9).

		Model number	% Correctly predicted (2010-2019 testing data)				% Correctly predicted (2020 testing data)			
			Energy poor	Non-energy poor	Accuracy	F1 score	Energy poor	Non-energy poor	Accuracy	F1 score
Random Forest:										
Max depth= 3	Estimators=1	1	64.9	69.3	67.1	66	65.8	69.5	67.7	65
“	Estimators=10	2	76.3	61.4	68.9	71	81.9	53.7	67.8	69
“	Estimators=50	3	79.4	59.6	69.5	72	85.3	52.0	68.7	70
Max. depth=6	Estimators=1	4	73.4	67.7	70.5	71	76.7	63.9	70.3	70
“	**Estimators=10**	**5**	**78.4**	**65.9**	**72.2**	**74**	**79**	**64.9**	**71.9**	**72**
“	Estimators=50	6	78.3	66.8	72.6	74	81.6	62.2	71.9	72
Max. depth=9	Estimators=1	7	80.6	63.3	71.9	74	77	65.3	71.2	71
“	Estimators=10	8	78	70	74	75	72.9	73.1	73	71
“	Estimators=50	9	77.7	70.8	74.3	75	74.1	70.6	72.3	71
Selected (boldfaced) Random Forest model with household-level variables only:										
Max. depth=6	Estimators=10	10	74.7	63.6	69.1	71	74.2	65.4	69.8	69
Decision Tree:										
Max depth= 3		11	64.5	71.7	68.1	67	62.2	73.7	67.9	64
Max depth= 6		12	77.9	65.1	71.5	73	74.5	67.7	71.1	70
Max depth= 9		13	76.5	71.2	73.8	75	70.6	74.4	72.5	70
k-Nearest Neighbors:										
Neighbors=2		14	60.6	82.1	71.4	85	35.7	83.9	59.8	46
Neighbors=5		15	75.7	69.9	72.8	80	59.5	71.7	65.6	62
Neighbors=10		16	69.6	75.7	72.6	76	54.9	77.6	66.2	61
XGBoost:										
Max. depth=6	Estimators=1	17	77.9	65.1	71.5	73	74.5	67.7	71.1	70
“	Estimators=10	18	76.6	71.5	74.1	75	69.3	75.3	72.3	70
“	Estimators=50	19	77.8	73.0	75.4	76	68.3	76.6	72.4	70

The performance measurement for the selected Random Forest model was further evaluated using an additional performance measurement, namely the Area Under the Curve (AUC) indicator (Bradley, 1997). The term “curve” stands for the Receiver Operating Characteristic Curve (ROC), i.e., the curve that plots model's TPR against its FPR at all classification thresholds (Bradley, 1997). In binary classification problems, AUC demonstrates the capability of a classifier to distinguish between classes; an AUC approaching 1 indicates strong capability, while an AUC approaching 0 indicates the opposite. Figure 2 demonstrates that the AUC of the selected Random Forest model was 0.78, indicating the model's strong capability to distinguish between the energy poor and non-energy poor classes.

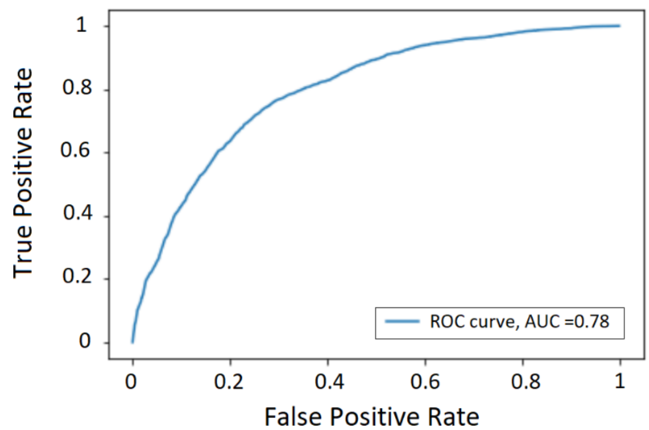

Figure 2: The selected Random Forest model's Receiver Operating Characteristic Curve (ROC). True and False Positive Rates are the proportions of positive classes that are classified correctly and incorrectly, respectively. The Area Under the Curve (AUC) value is 0.78.

4.2 Variable importance

Figure 3 illustrates the ranking of the selected Random Forest model's predictor variables in terms of their importance, based on the *PI* measure described in Section 3.3. Following common practice in the literature, the displayed values are normalized with the highest one being equal to 100 and the rest reported relative to the highest value (van Hove et al., 2022). The figure indicates that all predictors, both household-level and country-level, are associated with positive *PI* scores, indicating that they are important contributors to the model's accuracy. Not surprisingly, household income is the most important predictor; however, other household-level characteristics, such as the condition of the dwelling (damaged or not) and the household type rank high. In terms of country-level variables, social protection payments is the most influential one (and second most important in total, after income) in terms of *PI* score, followed by the gas supplier switching rates, the price of gas and the heating energy efficiency level achieved in households.

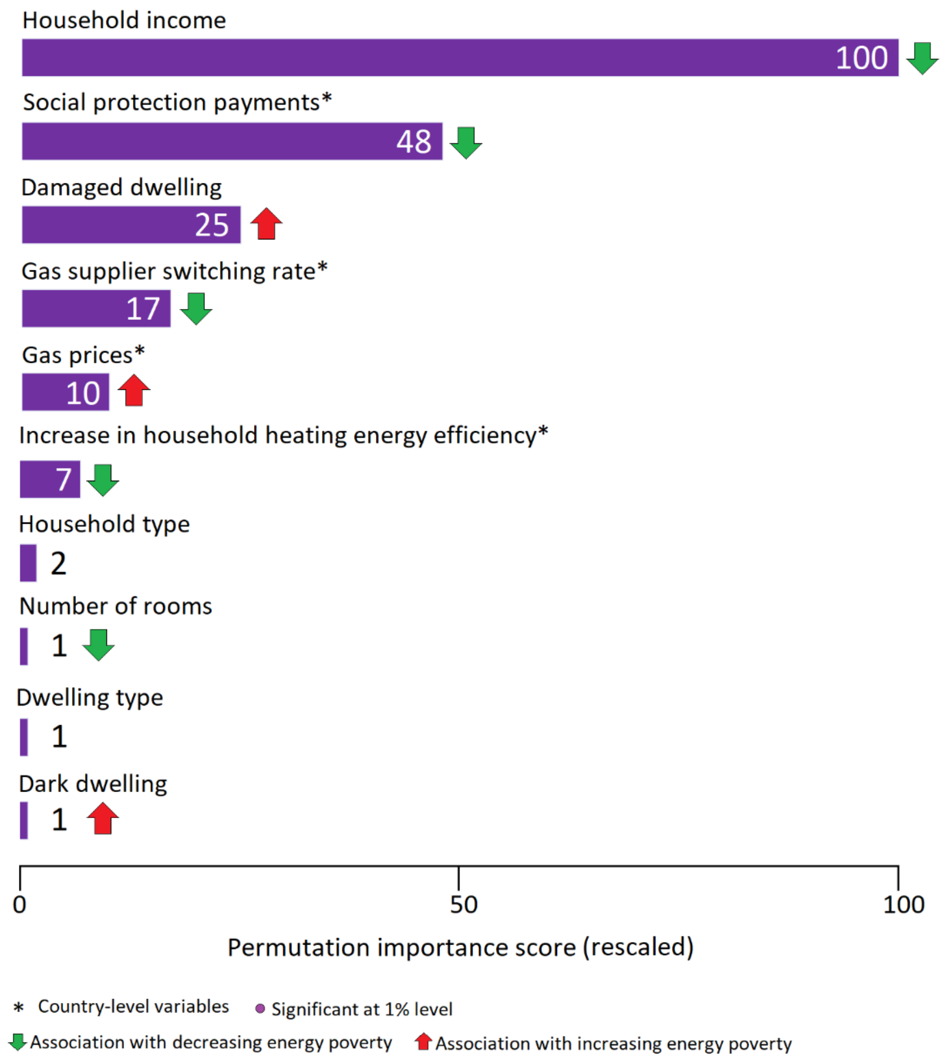

Figure 3: Relative importance of explanatory variables in predicting energy poverty. The importance is measured in contribution to the Random Forest model's accuracy, based on the Permutation Importance (*PI*) measure described in subsection 3.3. The figure additionally displays the statistical significance (purple color) and the direction of relationships (green and red arrows) among variables, as indicated by the logistic regression model. No directions are indicated for household type and dwelling type, as these may vary according to the particular types (see Table A.2)

While the *PI* measure ranks the predictors based on their contribution to the accuracy of the machine learning model, it does not reveal the direction of the relationships among variables and classes (positive or negative), nor their statistical significance. To investigate the statistical relevance of the explanatory variables, we include in our analysis a logistic regression model. Table A.2 of the

Appendix reports the resulting odds ratios, with an odds ratio greater than 1 signifying a positive association between predictors and outcome, and an odds ratio smaller than 1 signifying a negative association. The table indicates that the vast majority of relationships are statistically significant at the 1% level. Note that starting with household-level predictors only and subsequently adding country-level predictors does not lead to changes in the significance or sign of relationships, indicating the robustness of the model (Spandagos et al., 2022). Furthermore, noting that the regression provides information on association and not on causality, Table A.2 also indicates that increases in household income and received cash benefits, as well as increases in a country's social protection payments, heating energy efficiency, and gas supplier switching rates are associated with decreases in energy poverty. The same is true regarding dwellings that are not damaged, not dark and provide households with a higher number of rooms for living. On the other hand, increases in a country's gas prices are associated with higher levels of energy poverty. Moreover, it is evident that certain household types (such as single parent households and two adult households with three children or more) and certain dwelling types (such as apartments in buildings with ten or less dwellings) are more positively associated with energy poverty compared to other types.

Overall, our Random Forest and logistic regression models reveal that the condition of a dwelling is one of the most significant household-level predictors of energy poverty. Note that the link between dwelling condition and energy poverty has been highlighted mostly in relation to the dwelling's insulation and energy efficiency state (Mulder et al., 2023). In this work, however, we focus on dwelling condition in relation to the existence of damages such as leaking roof, damp walls or foundations, and rot in window frames or floors. Our findings contribute to achieving new understanding about the usefulness of this variable, which has been traditionally viewed as only a secondary and indirect predictor of energy poverty (Kyprianou et al., 2019). The connection between energy poverty and dwelling condition that excludes efficiency is most probably indirect indeed, as inability to fix damages in the dwelling (or to afford to live in a dwelling without them) plausibly accompanies inability to meet other needs, including the need for adequate heating. Nevertheless, the strength in the relationship between the two variables revealed here provides rationale for including dwelling condition in the main set of energy poverty predictors. As for the country-level variables that are revealed to exhibit high *PI* scores and be significantly associated with energy poverty reduction, (monetary) social protection and energy efficiency correspond to measures that are already of central importance within the European energy poverty alleviation strategies. On the other hand, facilitating

consumer switching of gas suppliers has not traditionally been a main priority in governmental agendas. However, the importance of the switching gas supply variable demonstrated here provides additional rationale for gradually increasing efforts to encourage consumers to perform such action.
We also used this model to simulate changes in energy poverty under different scenarios of policy interventions regarding gas prices, energy efficiency and social protection payments; the results of that exercise are illustrated in Figure A.1 of the Appendix. As that figure suggests, even the increase of country-level energy efficiency by 10%[1] has a smaller effect on reducing the probability of being energy poor, compared to increasing social protection payments alone. Overall, the observations derived from that figure exhibit the considerable effect of the most recent, real increases in gas prices and inflation on increasing the energy vulnerability of European consumers. It would be challenging to fully reverse this damage, even by jointly increasing social protection payments and energy efficiency. Therefore, additional ways to reverse the energy poverty effects of price and inflation increase should be actively sought after. The above highlight the urgency for developing effective targeting mechanisms to provide energy assistance to the energy vulnerable households.

4.3 Energy poverty alleviation targeting for equity and justice

The previous subsection highlighted the urgency for improving the targeting of energy poverty alleviation policies, enabling them to effectively provide energy assistance to energy vulnerable households. Such policies should contribute to the just character of the energy transition by being transparent, so that citizens can trust the authorities' decisions. Furthermore, it is crucial for such policies to increase fairness, or at least to minimize the unfairness that manifests when households that are indeed energy vulnerable do not receive energy assistance due to ambiguously chosen criteria. In this subsection, we discuss the usefulness of particular machine learning classifiers as tools to achieve these objectives.

4.3.1 Transparency and accountability

[1] Which is neither easy nor quick a task: the 2010-2020 Odyssee-Mure data reveals that it often takes more than 10 years for a EU-28 country to achieve such increases.

A comprehensive energy justice framework should include (among others) transparency and accountability in decision-making (Carley and Konisky, 2020). While discussions about procedural energy justice focus predominantly on who is included in decision-making processes, transparency and accountability pertaining to who (and why) is receiving energy poverty assistance or not can further emphasize equity and justice as key components of the energy transition. Any assistance scheme will result in people being excluded, therefore its underlying assigning mechanism should be transparently communicated to the general public to increase trust and the sentiment of fairness. To that end, employing algorithms that offer clear visualization of assigning mechanisms may facilitate the communication of targeting scheme details to the public. Among popular machine learning algorithms, Decision Trees are deemed particularly appropriate for such purposes, due to their ability to provide straightforward easily interpretable visualizations of their assigning mechanism (Andini et al., 2018). Here, we demonstrate the application of a Decision Tree model on our dataset and discuss its contribution to transparency compared to our selected Random Forest model. To facilitate comparison, we employed a Decision Tree model with a maximum depth of 6 layers (indicated as model number 11 on Table 2), which is the same as the one of the Random Forest model.

The assignment mechanism of the Decision Tree model is illustrated in Figure 4. This visualization reveals which parameters (and their combinations) the model employs to determine which households are classified either as energy poor or as non-energy poor. Furthermore, it reveals which thresholds determine the classification. For instance, the assigning mechanism has household income on its root node and the threshold of €25,591. This is a useful starting point in determining which households should receive assistance; however, as straightforward interpretation of this visualization is that targeting based on the criterion of household income only (and the specific threshold) would not be effective in reaching all the energy poor households and not reaching any of the non-energy poor households. Indeed, the left-hand side of the root node includes only households with an income lower than (or equal to) the threshold, while the right-hand side includes only households with an income higher than the threshold. If all left-hand side leaf nodes were corresponding to energy poor households and all the right-hand side ones to non-energy poor households, it would be indicated that such a targeting criterion is perfectly effective. Cleary, this is not the case in Figure 5; instead, both classes appear on both sides of the root node, and it is a combination of income with the additional criteria of energy efficiency, social protection payments, gas prices, dwelling condition and household type that eventually determines the classification. For instance, households with income of equal to or less than

€25,591 can still be non-energy poor if they are situated in countries with certain levels of social protection payments and increases in energy efficiency. Based purely on income as a criterion to provide energy poverty assistance, those household would be recipients without being energy poor. On the contrary, households with income of more than €25,591 can still be energy poor if other conditions pertaining to the additional criteria are true. Following solely an income-based criterion for providing assistance, those household would be non-recipients even though they are energy poor.

Observing how each branch leads to the classification outcomes in Figure 4 further reveals that the majority of the model's inferences are plausible. For instance, the majority of leaf nodes situated immediately on the left-hand side (representing the "less than" direction) of nodes that contain the variables of income, social protection payments and increase in energy efficiency belong to the energy poor class. This trend is plausible, as it is reasonable to assume that households below each income, benefit or efficiency threshold are closer to energy poverty compared to the households above it. The same is true for leaf nodes that are situated immediately below nodes indicating that households have children or live in a dwelling that is dark or damaged. On the other hand, a limited number of observed inferences are counter intuitive. For instance, leaf nodes situated immediately on the "less than" direction of nodes containing the gas prices variable belong to the energy poor class, while the opposite would have been expected. Furthermore, on one occasion, the threshold of cash benefits has a negative sign. These observations are examples that justify why the accuracy of the Decision Tree model in determining the energy poor class is not higher than 74.5% (as Table 2 indicates), and less than the one of our selected Random Forest model.

Note that Random Forest algorithms can also provide visualizations of their underlying assigning mechanisms. Nevertheless, a single Random Forest model comprises several estimators, with each estimator having its own assignment mechanism, root nodes, leaf nodes and thresholds. In comparison to the previous Decision Tree visualization, Figure 5 illustrates the assigning mechanism of a typical estimator in our selected Random Forest classifier. The particular estimator is chosen here for illustration to facilitate comparison with the Decision Tree mechanism, as they both have household incomes on their root node and the same threshold of €25,591. However, it often occurs that not all estimators of the same Random Forest model have the same predictor variable on their root note. Figure 5 demonstrates that the Random Forest visualization offers the same insights as the Decision Tree one does, namely decision variables, decision thresholds and the ability to examine the plausibility of inferences. Furthermore, it exhibits a far higher number of plausible inferences compared to the Decision Tree mechanism. For instance, leaf nodes immediately on the right-hand side of nodes containing the gas prices variable belong to the energy poor class, reflecting the plausible influence of higher prices in increasing energy poverty. Moreover, no unexpected negative signs are observed. The only observed inference that is less intuitive pertains to a node containing the variable of gas supplier switching rates, with the "more than" direction leads to an energy poor leaf node, contrary to what would have been expected. A limited number of cases like that might also appear on the other estimators comprising the forest. That phenomenon, however, does not damage the overall accuracy of the model, as the power of Random Forest models stems from the fact that they comprise several estimators, so that a certain weakness of a single estimator is counterbalanced by the strength of the others. However, while this characteristic of Random Forest algorithms contributes to their higher accuracy, it hinders their usefulness in providing transparency. Direct interpretation based on one or more isolated Random Forest estimators is not possible and the modeler can draw useful insights only by consulting the collective outcome of all estimators. Therefore, employing Random Forests for transparency and accountability at least require presenting all the assigning mechanisms of all estimators, increasing the complexity of information that needs to be shared. Even then, however, it would not be directly clear to the viewer how the different mechanisms of the estimators interact with each other to shape an accurate prediction outcome. Therefore, when the main objective is transparency and not accuracy, Decision Tree algorithms are a more effective basis of designing an alleviation scheme, due to the simplicity of sharing only one assignment mechanism instead of multiple.

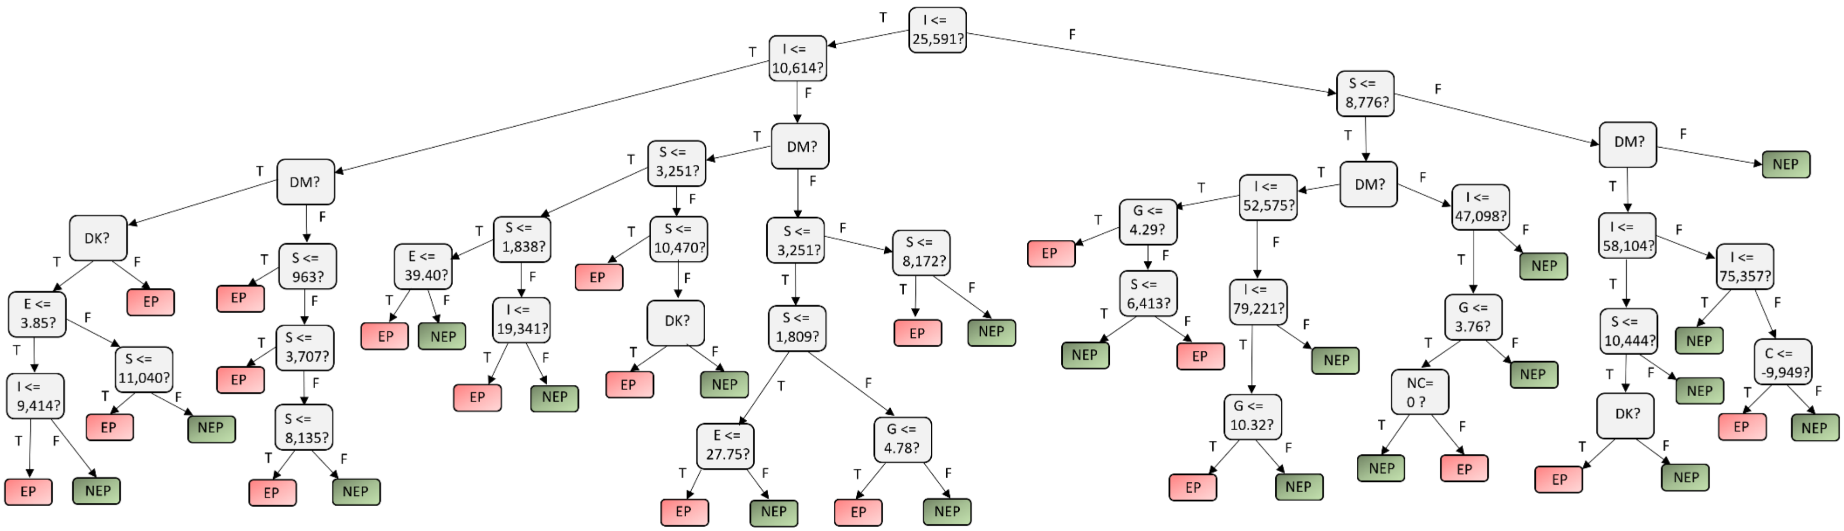

Figure 4: Illustration of the assigning mechanism of the selected Decision Tree model (model number 5 of Table 2). Colored nodes are leaf nodes. Abbreviations: EP, energy poor; NEP, non-energy poor; T, true; F, false, I, household income (€); C, cash benefits received (€); DM, damaged dwelling; DK, dark dwelling; NC, household type's number of children; S, country's social protection payments (€/capita); G, country's gas prices (€/GJ); E, country's increase in heating energy efficiency level for households since 2000 (%), SW, country's gas supplier switching rate (%)

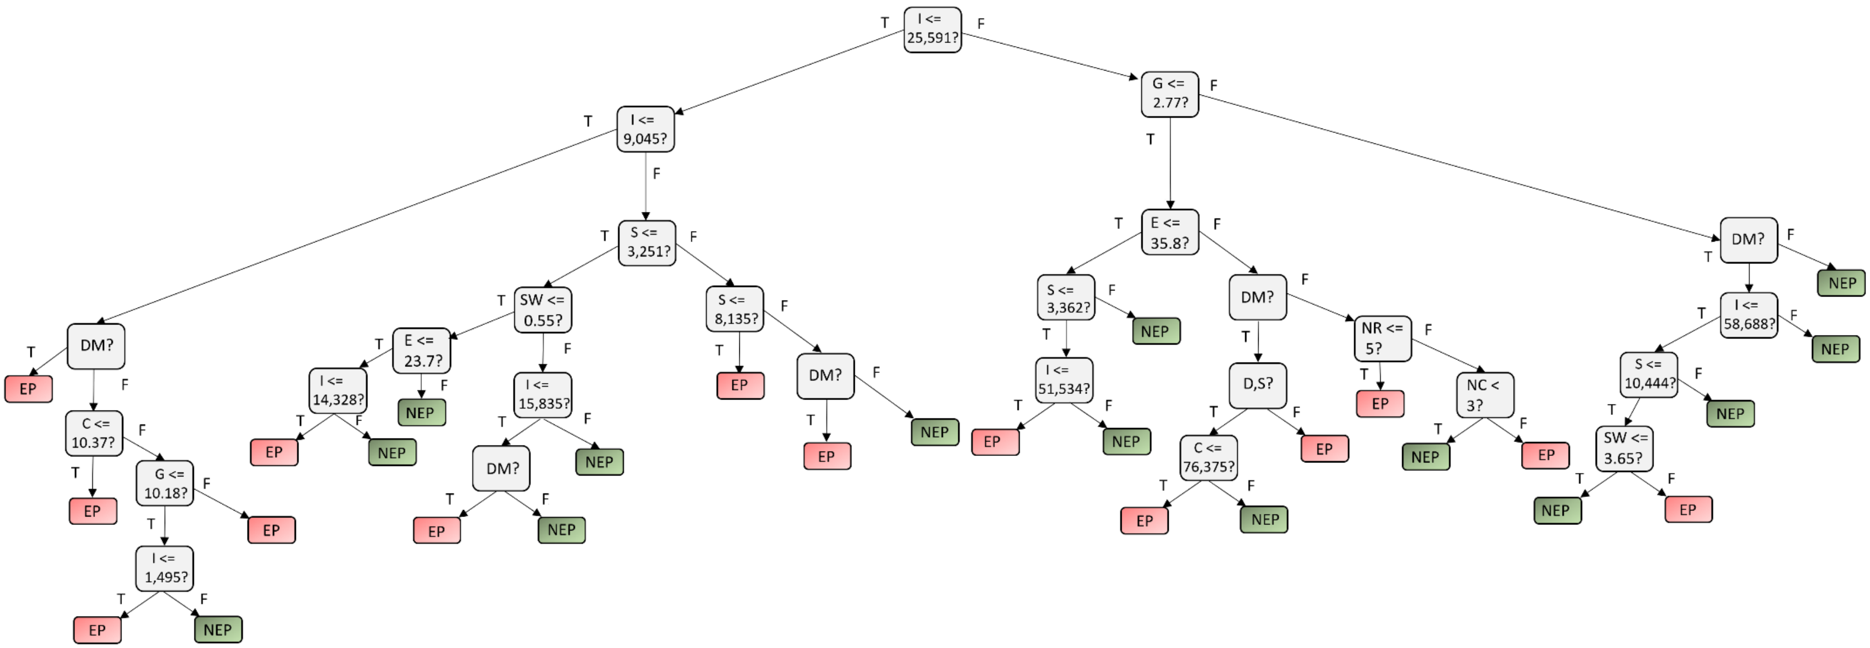

Figure 5: Illustration of the assigning mechanism of a typical estimator in the selected Random Forest model (model number 5 of Table 2). Colored nodes are leaf nodes. Abbreviations: EP, energy poor; NEP, non-energy poor; T, true; F, false, I, household income (€); C, cash benefits received (€); DM, damaged dwelling; D,S, detached or semi-detached house; NR, number of rooms; NC, household type's number of children; S, country's social protection payments (€/capita); G, country's gas prices (€/GJ); E, country's increase in heating energy efficiency level for households since 2000 (%), SW, country's gas supplier switching rate (%).

4.3.2 Fairness evaluation of income-relevant and social welfare-relevant targeting

The ability to properly identify those who need energy assistance is a prerequisite for developing fair energy poverty alleviation polices. In turn, alleviation policies that are fair can be an integral part of comprehensive policy packages focusing on the equity and justice aspects of the energy transition. In particular, identifying and recognizing who really needs energy poverty assistance can be a stepping stone towards distributional energy justice, where the objective is to ensure that no populations receive unfair shares of the transition's burdens (Carley and Konisky, 2020). Here, we conceptualize fairness in energy poverty targeting as the ability of alleviation schemes to properly identify energy poor households, with a particular focus on limiting the cases where households in need of energy assistance are excluded from it due to improperly set criteria. Contrary to the previous subsection, the main objective here was not to provide easily interpretable visualizations, but instead quantifiable results that are as accurate as possible. Therefore, we employed our selected Random Forest classifier instead of the Decision Tree one.

Building on the logic of identifying households that are energy poor and at the same time non-recipients of energy assistance, we estimated the fairness of hypothetical energy poverty alleviation schemes that are based on typical income-relevant or social welfare-relevant criteria. To avoid possible shortcomings associated with consulting the results of one estimator only, we estimated the schemes' fairness by consulting the collective outcome of all estimators in our Random Forest model. Using the 2020 data, we explored the following hypothetical scenarios:

1) <u>Alleviation schemes based on income-relevant criteria</u>

- Scenario 1A- Income thresholds: Only households below a particular income threshold will be recipients of energy assistance. We analyzed various hypothetical income thresholds in the €15,000-50,000 range to cover a wide range of realistic household incomes across the EU-28 countries. We identified which households in our 2020 sample would be energy assistance recipients based on each threshold; all other households were identified as non-recipients.
- Scenario 1B- 10% rule: Only households that spend more than 10% of their income on energy services will receive energy assistance.

- Scenario 1C- LIHC indicator: Only households whose energy expenditure exceeds the median in their country and their income minus that expenditure falls below the country's poverty line will receive energy assistance.

Scenarios 1B and 1C consider two of the most common objective and income-relevant energy poverty metrics, namely the 10% rule and the LIHC indicator. Running these scenarios requires information on household energy expenditure, which was not readily available in the EU-SILC databases. To overcome this issue, we estimated a regression model using common variables (such as income, household type and country fixed effects) from the Household Budget Survey (Eurostat, 2023c) database. We combined the estimated regression coefficients with the corresponding EU-SILC variables (income, household type and country fixed effects) to estimate energy expenditure levels for our 2020 sample. See Akoğuz et al. (2021) for more details about this procedure. Note that the imputed energy expenditure was only used to identify those households with potential high expenditure. Based on the energy expenditure estimations, we identified which households spend more than 10% of their income on energy. These would be the assistance recipient households under scenario 1B; all other households were identified as non-recipient. Similarly, we identified which households spend more than the median on energy while their income minus energy expenditures falls below their country's poverty line (60% of the country's median income- as provided by Eurostat). These would be the recipient households under scenario 1C, while all other households were identified as non-recipient.

2) <u>Alleviation schemes based on social welfare-relevant criteria</u>

- Scenario 2A- Old age: Only households having a member that is eligible for old age-based social payments will receive energy assistance.
- Scenario 2B- Unemployment: Only households having a member that eligible for unemployment-based social payments will receive energy assistance.

For this batch of scenarios, we estimate the fairness of hypothetical energy poverty alleviation schemes that would provide assistance based on two common social welfare-relevant criteria for householders, namely being classified as elderly and being unemployed. In our 2020 sample, we identified which households receive old age benefits or unemployment benefits in their countries; these would be the recipient households under scenarios 2A and 2B respectively, and all other households were identified as non-recipients.

For each of the scenarios, we determined which households within the pool of the non-recipient ones are classified as energy poor by our Random Forest model. This was achieved by combining information from the assigning mechanisms of all the estimators of the model. Table 3 summarizes the results for particular household categories. To facilitate result interpretation, we created three household categories by appropriately merging the household types. This resulted in 2 household categories with comparable representation in the 2020 testing dataset, namely “one person or single parent” with 13,251 households and “2 adults or more, no children”, with 14,454 households. The third category of households (“2 adults or more, children”) has a smaller representation with 7,571 households. The first block of columns on Table 3 provides the number of non-recipient households while the second block displays non-recipient households that are classified by our Random Forest model as energy poor. In practice, the provided percentage represents households that would be incorrectly identified as non-energy poor by the schemes and thus unfairly excluded from assistance. Therefore, the higher the percentage, the lower the scheme’s potential fairness. The percentages provided in brackets in the second block of columns are measured with respect to the total non-recipient households provided in the first block (in bold numbers).

Regarding Scenario 1A, overall, Table 3 indicates that for a variety of income thresholds, the hypothetical schemes would have excluded noteworthy percentages of households that are energy poor. Specifically, approximately 16-33% of the non-recipient households of all categories across the income thresholds are found to be in energy poverty. Among household categories, the highest percentages of the incorrectly excluded concern households consisting of 2 or more adults without children, and the second highest those consisting of 2 or more adults with children. Furthermore, Table 3 indicates that an alleviation scheme based exclusively on the 10% rule would result in approximately 44% of the non-recipient households being incorrectly excluded from assistance; the corresponding percentage under an alleviation scheme based exclusively on the LIHC indicator approximates 54%. In terms of social welfare-relevant criteria, alleviation schemes based exclusively on old age would result in approximately 51% of the non-recipient households being incorrectly excluded from assistance. Finally, the corresponding percentage under a scheme based exclusively on unemployment status approximates 56%. Under each one of the last four criteria, the highest percentages of the incorrectly excluded concern one person or single parent households- which have the second highest representation in our 2020 sample.

Table 3: Households (by category) that are non-recipient of hypothetical energy poverty alleviation schemes based on income-relevant and social welfare-relevant criteria, and fraction (boldfaced) of the non-recipients that are predicted to be energy poor by the selected Random Forest model- thus would had been incorrectly excluded from receiving energy assistance if the schemes were in force. The figures are based on the 2020 testing dataset.

EU-28, 2020 testing dataset (total: 35,276 households)								
Criteria for receiving energy assistance	Non-recipient households based on criterion/threshold				Non-recipient households that are also energy poor thus, incorrectly excluded from receiving assistance (% of all non-recipients)			
	One person or single parent	2 adults or more, no children	2 adults or more, children	**Total**	One person or single parent	2 adults or more, no children	2 adults or more, children	Total
Scenario 1A: Income thresholds								
-Below €15,000	4,279	9,298	6,251	**19,828**	1,444 (7.3%)	3,111 (15.7%)	2,028 (10.2%)	6,583 **(33.2%)**
-Below €20,000	3,027	7,376	5,432	**15,835**	714 (4.5%)	1,879 (11.9%)	1,496 (9.4%)	4,089 **(25.8%)**
-Below €25,000	2,200	5,999	4,655	**12,854**	365 (2.8%)	1,242 (9.7%)	1,130 (8.8%)	2,737 **(21.3%)**
-Below €30,000	1,743	5,002	4,019	**10,764**	247 (2.3%)	935 (8.7%)	897 (8.3%)	2,079 **(19.3%)**
-Below €35,000	1,320	4,229	3,504	**9,053**	174 (1.9%)	726 (8.0%)	720 (8.0%)	1,620 **(17.9%)**
-Below €50,000	620	2,687	2,400	**5,707**	73 (1.3%)	394 (6.9%)	434 (7.6%)	901 **(15.8%)**
Scenario 1B: 10% rule	7,731	9,507	6,143	**23,381**	4,703 (20.1%)	3,490 (14.9%)	2,031 (8.7%)	10,224 **(43.7%)**
Scenario 1C: LIHC indicator	13,221	12,284	6,929	**32,434**	9,197 (28.35%)	5,704 (17.6%)	2,634 (8.1%)	17,535 **(54.1%)**
Scenario 2A: Old age	6,636	5,427	6,626	**18,689**	4,341 (23.2%)	2,459 (13.2%)	2,695 (14.4%)	9,495 **(50.8%)**
Scenario 2B: Unemployment	12,375	13,087	6,396	**31,858**	8,657 (27.2%)	6,667 (20.9%)	2,550 (8.0%)	17,874 **(56.1%)**

5. Conclusions and policy implications

It is now timelier than ever to provide scientific guidance to policy-makers on improving the targeting of energy poverty schemes. Firstly, we are experiencing a global energy crisis that was unfolded during 2021 and escalated in 2022 (von Homeyer et al., 2022), creating supply fears and skyrocketing prices (Steckel et al., 2022) in Europe and elsewhere. Secondly, against the background of rising income and energy-relevant inequalities, several governments are planning to boost their carbon-neutrality efforts by implementing or increasing carbon taxation, and it is argued that such schemes require the parallel implementation of mechanisms that recycle the profits back to the households. For instance, Ireland's Climate Action Plan 2021 foresaw successive carbon tax increases to €100/ton by 2030 and a recycling of the profits back to households in the form of social benefits and energy efficiency assistance. To that end, it becomes a major policy challenge to design a recycling mechanism that is equitable and just, and the ability to correctly identify energy vulnerable households will be crucial in defining the equity outcomes of the energy transition. Finally, effective targeting may become critical for the success of EU programs such as the Just Transition Fund, a part of the European Green Deal, if funding is spent towards providing energy assistance.

In response to these challenges, we have successfully applied machine learning to predict energy poor households from a large and heterogeneous data pool comprised 28 European countries. Having the ability to satisfy different objectives, the resulting modeling framework was able to make accurate predictions or provide straightforward and transparent visualizations of the rules and thresholds that determine the classification between energy poor and non-energy poor households. In so doing, we provide a considerable advance on the literature by mapping from objective data to subjective self-reported inability to heat homes. The predictions with our selected Random Forest model were more accurate compared to those deriving from k-Nearest Neighbors algorithms; furthermore, they were comparable to those deriving from XGBoost algorithms, yet more accurate in predicting the class of main focus for this study, i.e., the energy poor class based on the latest available (2020) testing data. Moreover, our selected Random Forest comprising both household-level and country-level data was more accurate compared to a model consisting of the former type of data only. Our approach to merge household-level with country-level predictors not only improved the accuracy of the Random Forest classification, but also contributed to the identification of predictors of both types that are universal in their association with energy poverty, i.e., valid across all the 28 European countries examined. Furthermore, the ability to provide predictions of comparably high accuracy for each one of the 2010-2019 and 2020 periods contributes to a modeling solution that can be potentially effective with new and previously unseen by the system datasets, without self-reported energy poverty data, and without the need for re-training.

Our empirical findings have several policy implications for energy poverty alleviation in the EU-28 group. As mentioned earlier in this paper, current alleviation efforts of member states stand on two pillars: financial instruments to support low-income households and investment in dwelling energy efficiency. Our findings reinforce the importance of these two measures: both are reflected by two country-level variables that rank high in terms of prediction importance in our Random Forest model, namely social protection payments and heating energy efficiency for households. Furthermore, our regression analysis confirms that higher values of these variables are associated with statistically significant reductions in energy poverty across Europe. Our simulations show that while continuing governmental efforts in assisting households with both financial interventions and energy efficiency improvements is essential for addressing energy poverty, additional ways to fully reverse the adverse effects of the most recent price and inflation increases should be actively sought after.

A possible additional measure that should receive increased attention is relevant to another country-level variable that ranks high in terms of prediction importance in our Random Forest model, namely the gas supplier switching rate. Recently, such a parameter was recognized as an important dimension of a wider structural vulnerability index that was used to investigate energy poverty in Europe (Recalde et al., 2019). However, it has received less attention in the extant literature compared, for instance, to energy efficiency. High levels of consumer switching rates are deemed a desirable characteristic of well-functioning energy markets, as they reflect consumer ability to choose from a wide variety of available options (Harold et al., 2019). However, such rates remain small in Europe. Recently, two surveys conducted by the Office of Gas and Electricity Markets revealed that energy poor consumers are less likely to switch suppliers in the UK (Ambrosio-Albala et al., 2020); our machine learning and logistic regression models' results offer additional evidence of this relationship, with high switching rates in particular being associated with a statistically significant decrease in energy poverty across Europe. From a policy perspective, this provides an additional rationale for formulating regulatory processes and practices that facilitate supplier switching. Currently, switching energy suppliers relies predominately on the ability and willingness of consumers to take initiative (Mahoney et al., 2020), but it is less common for low-income households to do so, due to energy debts (Middlemiss, 2017), uncertainty or lack of information (Lorenc et al., 2013). In addition, poorly-educated consumers might also be less able to switch, due to lower self-efficacy during negotiations (Sheehy-Skeffington and Rea, 2017). An important step has been taken by certain European countries, such as Denmark, France, Luxemburg and the UK, by allowing supplier switching to consumers even when indebted (Pye et al., 2015a). However, supplier switching needs to be further facilitated through information provision campaigns and programs that provide individual assistance to households that are less able to identify and negotiate more favorable terms. To that end, continued research (such as Ambrosio-Albala et al.

(2020), for instance) on understanding how energy poor households interact with the energy supply market should be conducted.

We have also demonstrated how machine learning classifiers can contribute to the pursuit of energy equity and justice by improving the transparency and evaluating the fairness potential of hypothetical energy poverty alleviation schemes. In this way, we make several contributions to overcoming targeting challenges. Firstly, we suggest the use of Decision Trees algorithms when the policy-makers' main objective is transparency and accountability pertaining to the criteria and thresholds that determine who will be a recipient of assistance. The Decision Tree classifier presented in this paper constitutes a transparent framework to visualize and understand the decision thresholds that define the energy poverty status classification across the EU-28, overcoming the "black box" limitation that often manifests in similar applications. This transparency is particularly important in creating public trust towards net-zero policies and encouraging citizen participation in the energy transition. Secondly, we suggest Random Forest algorithms when the policy-makers' main objective is higher accuracy in identifying who should be recipient of energy assistance- and ensuring that no energy poor households are unfairly excluded from it. Our selected Random Forest classifier suggests that social welfare systems and their typical criteria should not be the only channels for identifying recipients of energy poverty programs. If that happens, a significant number of energy poor households in need of assistance might be excluded from it, limiting the fair character of schemes. In particular, schemes based on income thresholds can lead to exclusion that is more noteworthy for households consisting of 2 adults or more and children. Furthermore, more noteworthy exclusion is observed for one person or single parent households when it comes to schemes based on the 10% rule, the LIHC indicator, householders' old age status or unemployment status. Therefore, the process of defining criteria for energy assistance provision should include additional characteristics to more accurately determine the energy vulnerability status. For instance, as the condition of a household's dwelling is revealed here as an important predictor of energy poverty, authorities may better identify the energy vulnerable by incorporating in their decision-making tools transparent information about the building stock of their jurisdictions. Apart from the EU-SILC surveys, platforms such as the EU Building Stock Observatory (which already incorporates energy poverty-relevant indicators) can be useful tools for monitoring and understanding the building condition challenges that European households face.

Finally, our inclusion of country-level variables in a machine learning framework evaluating energy poverty alleviation schemes is a new departure that serves the purpose of providing an EU-wide understanding of key trends. This is particularly necessary in the absence of a standardized framework to compare energy poverty and relevant trends across European countries of different social policy and energy market characteristics. It will be up to the policy-makers to decide whether the context they operate in requires tools that favor transparency or accuracy, or a combination of both. Nevertheless, combining the ability to offer transparency with the inclusion of predictors from two levels may serve two additional purposes: the transparency offered pertaining to household-level predictors can help policy-makers to set more realistic thresholds for granting assistance to households; at the same time, the transparency offered pertaining to country-level predictors may assist EU-level decisions towards allocating assistance, for instance through the Just Transition Fund or similar programs, to countries that fall behind the indicated thresholds in social protection payments, energy efficiency and supplier switching rates.

The results of this paper should be considered in light of certain limitations. The variable representing energy poverty that was included in the model training process derived from answers to the EU-SILC survey, which concerned perceived energy poverty- only one of the possible manifestations of the phenomenon. Moreover, cognitive bias and error of exclusion might exist in survey answers, as well as desirability bias, especially given the sensitive nature of commenting on the ability to keep one’s home adequately warm. However, survey-based energy poverty indicators in general, remain the preferred choice among European policy-makers (Karpinska and Śmiech, 2021), especially when the target is to measure vaguely-defined concepts. Furthermore, anonymity is guaranteed in the EU-SILC methodology to limit biases as much as possible. Most importantly, it has been demonstrated that survey respondents who identified themselves (or their households) as energy poor are more likely to be objectively energy poor (Cong et al., 2022). Nevertheless, recognizing the need to further overcome such limitations in the future, the framework presented here contributes to the development of modeling solutions that will be gradually becoming less dependent on self-reported data.

The machine learning approaches presented here may be expanded in the future to provide a deeper understanding of the various dimensions of alleviating energy poverty within the energy transition. We suggest four possible, data-dependent directions for such an expansion: i) Incorporating data that provide insights on supplementary factors that may have shaped certain key variables regarding energy poverty prediction. For instance, supplier switching rates are demonstrated here to be an important predictor, and it would be intriguing to examine data on underlying processes that may have facilitated or barred this practice, for instance whether households have negotiated with their suppliers to get favorable terms or used energy price comparison services. ii) Expanding the geographical scope. While the data and policy implications discussed here were EU-centric, the approach we demonstrated may be used to derive equally meaningful insights for other jurisdictions, if fed with the appropriate data. Independently of the definitions or metrics used in each jurisdiction, it is evident that energy poverty is becoming prominent both in developed and developing countries. For instance, it is estimated that 31% of households in the US face some form of energy poverty (U.S. Energy Information Administration (EIA)), while corresponding rates in other populous countries such as China and India have been estimated to be 18.9% (Lin and Wang, 2020) and 65% (Gupta et al., 2020), respectively. Therefore, continuous research with expanded geographical focus is needed to support alleviation efforts at the global level. To that end, integrating community-level data (Huang et al., 2022) with household-level and country-level data in any examined country may further contribute to the development of more accurate simulations. Note that the current work did not use community-level data in the context of all the EU countries, as such data were not available in the EU-SILC databases. However, future work may benefit from such data once they become available. iii) Incorporating novel metrics that represent additional behavioral patterns of consumers and provide additional insights on the equity and justice dimensions of the energy transition. An example of such newly-proposed metrics is the "energy equity gap" (Cong et al., 2022), which concerns the outdoor temperature at which households of various income levels start using their cooling systems. Focusing on consumer behaviors relevant to cooling will be increasingly important in future energy poverty research, as income and population increase in the world's hottest countries is expected to make cooling energy one of the main drivers of global energy demand (International Energy Agency (IEA), 2018). iv) Finally, the possibility of using machine learning models based on a multidimensional energy poverty approach would be a fruitful ground for future research. This would require data on energy use for end uses such as heating water or cooking food (as employed in Huang et al. (2022) and Nussbaumer et al. (2012)), which are currently not available through the EU-SILC surveys.

The availability of big sets of relevant data will be an essential determinant of success for work in these four directions. Nevertheless, Random Forest and other machine learning tools are proven capable of handling big datasets and transforming them into meaningful insights and actionable recommendations for energy poverty alleviation at the international scale.

REFERENCES

Abbas, K., Manzoor, K., Xu, D., Ali, M., Baz, K., Hussain, S., Ahmed, M., 2022. Measurements and determinants of extreme multidimensional energy poverty using machine learning. Energy 251, 123977. https://doi.org/10.1016/j.energy.2022.123977

Akoğuz, E.C., Capéau, B., Decoster, A., De Sadeleer, L., Güner, D., Manios, K., Paulus, A., Vanheukelom, T., 2021. A new indirect tax tool for EUROMOD Final Report (No. Final Report JRC Project no. JRC/SVQ/2018/B.2/0021/OC).

Ambrosio-Albala, P., Middlemiss, L., Owen, A., Hargreaves, T., Emmel, N., Gilbertson, J., Tod, A., Snell, C., Mullen, C., Longhurst, N., Gillard, R., 2020. From rational to relational: How energy poor households engage with the British retail energy market. Energy Research & Social Science 70. https://doi.org/10.1016/j.erss.2020.101765

Andini, M., Ciani, E., de Blasio, G., D'Ignazio, A., Salvestrini, V., 2018. Targeting with machine learning: An application to a tax rebate program in Italy. Journal of Economic Behavior and Organization 156, 86–102. https://doi.org/10.1016/j.jebo.2018.09.010

Bednar, D.J., Reames, T.G., 2020. Recognition of and response to energy poverty in the United States. Nature Energy 5. https://doi.org/10.1038/s41560-020-0582-0

Beiser-McGrath, L., Huber, R., 2018. Assessing the relative importance of psychological and demographic factors for predicting climate and environmental attitudes. Climatic Change 149, 335–347.

Best, R., Hammerle, M., Mukhopadhaya, P., Silber, J., 2021. Targeting household energy assistance. Energy Economics 99, 105311. https://doi.org/10.1016/j.eneco.2021.105311

Biau, G., Scornet, E., 2016. A Random Forest Guided Tour. TEST 25, 197–227.

Bienvenido-Huertas, D., Pulido-Arcas, J., Rubio-Bellido, C., Pérez-Fargallo, A., 2021. Prediction of Fuel Poverty Potential Risk Index Using Six Regression Algorithms: A Case-Study of Chilean Social Dwellings. Sustainability 13.

Boardman, B., 2009. Fixing Fuel Poverty: Challenges and Solutions. Routledge, London. https://doi.org/https://doi.org/10.4324/9781849774482

Bouzarovski, S., Petrova, S., 2015. A global perspective on domestic energy deprivation: Overcoming the energy poverty- fuel poverty binary. Energy Research & Social Science 10, 31–40.

Bradley, A.E., 1997. The use of the area under the ROC curve in the evaluation of machine learning algorithms. Pattern Recognition 30, 1145–1159.

Breiman, L., 2001. Random Forests. Machine Learning 45, 5–32.

Carley, S., Konisky, D., 2020. The justice and equity implications of the clean energy transition. Nature Energy 5, 569–577.

Chakrabarty, N., Biswas, S., 2018. A Statistical Approach to Adult Census Income Level Prediction, in: 2018 International Conference on Advances in Computing, Communication Control and Networking (ICACCCN).

Cong, S., Nock, D., Qiu, Y., Xing, B., 2022. Unveiling hidden energy poverty using the energy equity gap. Nature Communications 13, 2456. https://doi.org/10.1038/s41467-022-30146-5

Curbelo Montañez, C., Hurst, W., 2020. A machine learning approach for detecting unemployment using the smart metering infrastructure. IEEE Access 8. https://doi.org/10.1109/ACCESS.2020.2969468

Dalla Longa, F., Sweerts, B., van der Zwaan, B., 2021. Exploring the complex origins of energy poverty in The Netherlands with machine learning. Energy Policy 156, 112373. https://doi.org/10.1016/j.enpol.2021.112373

Department of Business Energy & Industrial Strategy (BEIS), 2017. Machine learning and fuel poverty targeting- Annex A.

Dogan, E., Madaleno, M., Inglesi-Lotz, R., Taskin, D., 2022. Race and energy poverty: Evidence from African-American households. Energy Economics 108. https://doi.org/10.1016/j.eneco.2022.105908

Dubois, U., 2012. From targeting to implementation: The role of identification. Energy Policy 49, 107–115. https://doi.org/10.1016/j.enpol.2011.11.087

Eurostat, 2023a. Energy consumption in households [WWW Document]. URL https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Energy_consumption_in_households

Eurostat, 2023b. EU statistics on income and living conditions [WWW Document]. URL https://ec.europa.eu/eurostat/web/microdata/european-union-statistics-on-income-and-living-conditions

Eurostat, 2023c. Household Budget Surveys [WWW Document]. URL https://ec.europa.eu/eurostat/web/household-budget-surveys

Feldmeyer, D., Meisch, C., Sauter, H., Birkmann, J., 2020. Using OpenStreetMap Data and Machine Learning to Generate Socio-Economic Indicators. International Journal of Geo-Information 9, 498.

Friedman, J., 2001. Greedy Function Approximation: A Gradient Boosting Machine. The Annals of Statistics 29, 1189–1232.

González-Eguino, M., 2015. Energy poverty: An overview. Renewable and Sustainable Energy Reviews 47, 377–385. https://doi.org/10.1016/j.rser.2015.03.013

Gupta, R., Gregg, M., 2018. Targeting and modelling urban energy retrofits using a city-scale energy mapping approach. Journal of Cleaner Production 174.

Gupta, S., Gupta, E., Sarangi, G.K., 2020. Household Energy Poverty Index for India: An analysis of inter-state differences. Energy Policy 144, 111592. https://doi.org/10.1016/j.enpol.2020.111592

Harold, J., Cullinan, J., Lyons, S., 2019. Consumer switching in European retail markets. Oxford Economic Papers 1–19. https://doi.org/10.1093/oep/gpz044

Hills, J., 2012. Getting the Measure of Fuel Poverty: Final Report of the Fuel Poverty Review CASE report 72.

Hong, Z., Park, I., 2021. Comparative analysis of energy poverty prediction models using machine learning algorithms. Journal of Korea Planning Association 56.

Hu, L., He, S., Han, Z., Xiao, H., Su, S., Weng, M., Cai, Z., 2019. Monitoring housing rental prices based on social media: An integrated approach of machine-learning algorithms and hedonic modeling to inform equitable housing policies. Land Use Policy 82, 657–673. https://doi.org/10.1016/j.landusepol.2018.12.030

Huang, Y., Jiao, W., Wang, K., Li, E., Yan, Y., Chen, J., Guo, X., 2022. Examining the multidimensional energy poverty trap and its determinants: An empirical analysis at household and community levels in six provinces of China. Energy Policy 169, 113193. https://doi.org/10.1016/j.enpol.2022.113193

International Energy Agency (IEA), 2018. The Future of Cooling. Opportunities for energy-efficient air conditioning.

Jiao, Y., Du, P., 2016. Performance measures in evaluating machine learning based bioinformatics predictors for classifications. Quantitative Biology 4, 320–330.

https://doi.org/10.1007/s40484-016-0081-2

Johns, N.E., Cowling, K., Gakidou, E., 2013. The wealth (and health) of nations: a cross-country analysis of the relation between wealth and inequality in disease burden estimation. The Lancet 381, S66. https://doi.org/10.1016/S0140-6736(13)61320-3

Karpinska, L., Śmiech, S., 2021. Breaking the cycle of energy poverty. Will Poland make it? Energy Economics 94, 105063. https://doi.org/10.1016/j.eneco.2020.105063

Kuncheva, L., 1995. Editing for the k-nearest neighbors rule by a genetic algorithm. Pattern Recognition Letters 16, 809–814.

Kyprianou, I., Serghides, D.K., Varo, A., Gouveia, J.P., Kopeva, D., Murauskaite, L., 2019. Energy poverty policies and measures in 5 EU countries: A comparative study. Energy & Buildings 196, 46–60. https://doi.org/10.1016/j.enbuild.2019.05.003

Le, T.T., Fu, W., Moore, J.H., 2020. Scaling tree-based automated machine learning to biomedical big data with a feature set selector. Bioniformatics 36, 250–256. https://doi.org/10.1093/bioinformatics/btz470

Lee, T., Markowitz, E., Howe, P., Ko, C., Leiserowitz, A., 2015. Predictors of public climate change awareness and risk perception around the world. Nature Climate Change 5.

Levi, S., 2021. Why hate carbon taxes? Machine learning evidence on the roles of personal responsibility, trust, revenue recycling, and other factors across 23 European countries. Energy Research & Social Science 73, 101883.

Li, K., Lloyd, B., Liang, X., Wei, Y., 2014. Energy poor or fuel poor: What are the differences? Energy Policy 68, 476–481. https://doi.org/10.1016/j.enpol.2013.11.012

Lin, B., Wang, Y., 2020. Does energy poverty really exist in China? From the perspective of residential electricity consumption. Energy Policy 143, 111557. https://doi.org/10.1016/j.enpol.2020.111557

López-Vargas, A., Ledezma-Espino, A., Sanchis-de-Miguel, A., 2022. Methods , data sources and applications of the Artificial Intelligence in the Energy Poverty context: A review. Energy & Buildings 268. https://doi.org/10.1016/j.enbuild.2022.112233

Lorenc, A., Pedro, L., Badesha, B., Dize, C., Fernow, I., Dias, L., 2013. Tackling fuel poverty through facilitating energy tariff switching: a participatory action research study in vulnerable groups. Public Health 127, 894–901. https://doi.org/10.1016/j.puhe.2013.07.004

Ma, J., Cheng, J.C.P., 2016. Identifying the influential features on the regional energy use intensity of residential buildings based on Random Forests. Applied Energy 183, 193–201.

https://doi.org/10.1016/j.apenergy.2016.08.096

Mahoney, K., Gouveia, J.P., Palma, P., 2020. (Dis) United Kingdom? Potential for a common approach to energy poverty assessment. Energy Research & Social Science 70. https://doi.org/10.1016/j.erss.2020.101671

Marcucci, A., Fragkos, P., 2015. Drivers of regional decarbonization through 2100: A multi-model decomposition analysis. Energy Economics 51, 111–124. https://doi.org/10.1016/j.eneco.2015.06.009

Middlemiss, L., 2017. A critical analysis of the new politics of fuel poverty in England. Critical Social Policy 37, 425–443. https://doi.org/10.1177/0261018316674851

Mulder, P., Dalla Longa, F., Straver, K., 2023. Energy poverty in the Netherlands at the national and local level: A multi-dimensional spatial analysis. Energy Research & Social Science 96. https://doi.org/10.1016/j.erss.2022.102892

Nobre, J., Neves, R.F., 2019. Combining Principal Component Analysis , Discrete Wavelet Transform and XGBoost to trade in the financial markets. Expert Systems With Applications 125, 181–194. https://doi.org/10.1016/j.eswa.2019.01.083

Nussbaumer, P., Bazilian, M., Modi, V., 2012. Measuring energy poverty: Focusing on what matters. Renewable and Sustainable Energy Reviews 16, 231–243. https://doi.org/10.1016/j.rser.2011.07.150

Nussbaumer, P., Nerini, F.F., Onyeji, I., Howells, M., 2013. Global Insights Based on the Multidimensional Energy Poverty Index (MEPI). Sustainability 5, 2060–2076. https://doi.org/10.3390/su5052060

Okushima, S., 2016. Measuring energy poverty in Japan, 2004 – 2013. Energy Policy 98, 557–564. https://doi.org/10.1016/j.enpol.2016.09.005

Ortega-Ruiz, G., Mena-Nieto, A., Golpe, A., García-Ramos, J., 2022. CO2 emissions and causal relationships in the six largest world emitters. Renewable and Sustainable Energy Reviews 162, 112435.

Pachauri, S., Mueller, A., Kemmler, A., Spreng, D., 2004. On Measuring Energy Poverty in Indian Households. World Development 32, 2083–2104. https://doi.org/10.1016/j.worlddev.2004.08.005

Pye, S., Dobbins, A., Ba, C., Brajkovi, J., Deane, P., De Miglio, R., 2015a. Addressing Energy Poverty and Vulnerable Consumers in the Energy Sector Across the EU. https://doi.org/10.3917/eufor.378.0064

Pye, S., Dobbins, A., Baffert, C., Brajković, J., Grgurev, I., De Miglio, R., Deane, P., 2015b. Energy poverty and vulnerable consumers in the energy sector across the EU: analysis of policies and measures.

Raffio, G., Isambert, O., Mertz, G., Schreier, C., Kissock, K., 2007. Targeting Residential Energy Assistance, in: ASME 2007 Energy Sustainability Conference. Long Beach, CA, pp. 489–496.

Reades, J., De Souza, J., Hubbard, P., 2019. Understanding urban gentrification through machine learning. Urban Studies 56.

Reames, T., 2016. Targeting energy justice: Exploring spatial, racial/ethnic and socioeconomic disparities in urban residential heating energy efficiency. Energy Policy 97, 549–558.

Recalde, M., Peralta, A., Oliveras, L., Tirado-Herrero, S., Borrell, C., Palència, L., Gotsens, M., Artazcoz, L., Marí-Dell'Olmo, M., 2019. Structural energy poverty vulnerability and excess winter mortality in the European Union: Exploring the association between structural determinants and health. Energy Policy 133, 1–18. https://doi.org/10.1016/j.enpol.2019.07.005

Roberts, D., Vera-Toscano, E., Phimister, E., 2015. Fuel poverty in the UK: Is there a difference between rural and urban areas? Energy Policy 87, 216–223.

Robinson, C., Bouzarovski, S., Lindley, S., 2018. Getting the measure of fuel poverty: The geography of fuel poverty indicators in England. Energy Research & Social Science 36, 79–93. https://doi.org/10.1016/j.erss.2017.09.035

Rogelj, J., den Elzen, M., Höhne, N., Fransen, T., Fekete, H., Winkler, H., Schaeffer, R., Sha, F., Riahi, K., Meinshausen, M., 2016. Paris Agreement climate proposals need a boost to keep warming well below 2 °C. Nature 534, 631–639. https://doi.org/10.1038/nature18307

Samuel, A., 1959. Some Studies in Machine Learning Using the Game of Checkers. IBM Journal of Research and Development 3.

Sareen, S., Thomson, H., Tirado, S., Lippert, I., Lis, A., 2020. European energy poverty metrics: Scales, prospects and limits. Global Transitions 2, 26–36. https://doi.org/10.1016/j.glt.2020.01.003

Serrano, E., del Pozo-Jiménez, P., Suárez-Figueroa, M., González-Pachón, J., Bajo, J., Gómez-Pérez, A., 2018. Predicting the risk of suffering chronic social exclusion with machine learning, in: Distributed Computing and Artificial Intelligence, 14th International Conference. pp. 132–139. https://doi.org/10.1007/978-3-319-62410-5

Sharma, D.D., Singh, S.N., Lin, J., 2017. Identification and characterization of irregular consumptions of load data. Journal of Modern Power Systems and Clean Energy 5, 465–477. https://doi.org/10.1007/s40565-017-0268-1

Sheehy-Skeffington, J., Rea, J., 2017. How Poverty Affects People's Decision-Making Processes. York.

Siksnelyte-Butkiene, I., Streimikiene, D., Lekavicius, V., Balezentis, T., 2021. Energy poverty indicators: A systematic literature review and comprehensive analysis of integrity. Sustainable Cities and Society 67. https://doi.org/10.1016/j.scs.2021.102756

Sovacool, B.K., 2012. The political economy of energy poverty: A review of key challenges. Energy for Sustainable Development 16, 272–282. https://doi.org/10.1016/j.esd.2012.05.006

Spandagos, C., Baark, E., Ng, T., Masaru, Y., 2021. Social Influence and Economic Intervention Policies to Save Energy at Home: Critical Questions for the New Decade and Evidence from Air-condition Use. Renewable and Sustainable Energy Reviews 143, 110915.

Spandagos, C., Tovar Reaños, M., Lynch, M.Á., 2022. Public acceptance of sustainable energy innovations in the European Union: A multidimensional comparative framework for national policy. Journal of Cleaner Production 340, 130721. https://doi.org/10.1016/j.jclepro.2022.130721

Steckel, J.C., Missbach, L., Ohlendorf, N., Feindt, S., 2022. Effects of the energy price crisis on European households and policy options. Berlin.

Strobl, C., Malley, J., Tutz, G., 2009. An Introduction to Recursive Partitioning: Rationale , Application , and Characteristics of Classification and Regression Trees , Bagging , and Random Forests. Psychological Methods 14, 323–348. https://doi.org/10.1037/a0016973

Tovar, M.A., 2021. Fuel for poverty: A model for the relationship between income and fuel poverty . Evidence from Irish microdata. Energy Policy 156. https://doi.org/10.1016/j.enpol.2021.112444

U.S. Energy Information Administration (EIA), n.d. One in three U.S. households faced challenges in paying energy bills in 2015 [WWW Document]. URL https://www.eia.gov/consumption/residential/reports/2015/energybills/

Uniejewski, B., Marcjasz, G., Weron, R., 2019. Understanding intraday electricity markets: Variable selection and very short-term price forecasting using LASSO. International Journal of Forecasting 35, 1533–1547. https://doi.org/10.1016/j.ijforecast.2019.02.001

van Hove, W., Dalla, F., van der Zwaan, B., 2022. Identifying predictors for energy poverty in Europe using machine learning. Energy & Buildings 264, 112064. https://doi.org/10.1016/j.enbuild.2022.112064

von Homeyer, I., Oberthür, S., Dupont, C., 2022. Implementing the European Green Deal during the Evolving Energy Crisis. Journal of Common Market Studies 1–12. https://doi.org/10.1111/jcms.13397

Walker, R., Mckenzie, P., Liddell, C., Morris, C., 2012. Area-based targeting of fuel poverty in Northern Ireland: An evidenced-based approach. Applied Geography 34, 639–649. https://doi.org/10.1016/j.apgeog.2012.04.002

Wang, H., Maruejols, L., Yu, X., 2021. Predicting energy poverty with combinations of remote-sensing and socioeconomic survey data in India: Evidence from machine learning. Energy Economics 102, 105510. https://doi.org/10.1016/j.eneco.2021.105510

Wang, Y., Lin, B., 2022. Can energy poverty be alleviated by targeting the low income? Constructing a multidimensional energy poverty index in China. Applied Energy 321, 119374. https://doi.org/10.1016/j.apenergy.2022.119374

Wang, Zi, Yu, H., Luo, M., Wang, Zhe, Zhang, H., Jiao, Y., 2019. Predicting older people' s thermal sensation in building environment through a machine learning approach: Modelling , interpretation, and application. Building and Environment 161, 106231. https://doi.org/10.1016/j.buildenv.2019.106231

Yin, H., Zhou, K., 2022. Performance evaluation of China's photovoltaic poverty alleviation project using machine learning and satellite images. Utilities Policy 76, 101378. https://doi.org/10.1016/j.jup.2022.101378

Zhang, C., Cao, L., Romagnoli, A., 2018. On the feature engineering of building energy data mining. Sustainable Cities and Society 39, 508–518. https://doi.org/10.1016/j.scs.2018.02.016

Appendix

Table A.1: Accuracies and F1scores of the predictions made by our selected Random Forest model (max. depth = 6, number of estimators= 10) with different years' datasets.

Year	Accuracy	F1 score
2020	71.9	72
2019	71.7	70

2018	71.8	72
2017	71.2	73
2016	71.2	74

Table A.2: Summary of the results of the logistic regression described in Section 4.2. Odds ratios lower than 1 and significant are boldfaced.

Energy poverty, EU-28				
	Household-level only model		**Household-level & country-level model**	
	Odds ratio		Odds ratio	
Household-level predictors				
Income (log)	**0.45*** **	**(0.00)**	**0.46*** **	**(0.00)**
Cash benefits received	**0.99*** **	**(0.00)**	**0.99*** **	**(0.00)**
Number of rooms	**0.89*** **	**(0.00)**	**0.89*** **	**(0.00)**
Dwelling not damaged	**0.44*** **	**(0.00)**	**0.44*** **	**(0.00)**
Dwelling not too dark	**0.60*** **	**(0.00)**	**0.60*** **	**(0.00)**
<u>Dwelling type:</u>				
-Semi-detached house	1.05***	(0.00)	1.04***	(0.00)
-Apartment in a building with 10 dwellings or less	1.09***	(0.00)	1.09***	(0.00)
-Apartment in a building with 10 dwellings or more	1.07***	(0.00)	1.07***	(0.00)
<u>Household type:</u>				
-2 adults under 65 y.o., no children	1.16***	(0.01)	1.14***	(0.01)
-2 adults, at least one older than 65 y.o., no children	1.02**	(0.00)	1.01	(0.00)
-Single parent, 1 child or more	1.56***	(0.02)	1.54***	(0.02)
-2 adults, 1 child	1.13***	(0.01)	1.10***	(0.01)
-2 adults, 2 children	1.13***	(0.01)	1.11***	(0.01)
-2 adults, 3 children or more	1.67***	(0.02)	1.63***	(0.02)
-Other households without children	1.73***	(0.01)	1.69***	(0.01)
-Other households with children	2.03***	(0.02)	1.97***	(0.02)
Country-level predictors				
Gas prices			1.02***	(0.00)
People having switched gas providers			**0.98*** **	**(0.00)**
Social protection payments per capita			**0.99*** **	**(0.00)**
Heating energy efficiency in households (increase since 2000)			**0.99*** **	**(0.00)**
Summary				
Cons	246.55	12.64	4466.53	856.30
Log pseudo likelihood	-676457.31		-675462.95	
Pseudo R^2	0.1982		0.1994	

Observations	2,419,500	2,419,500
(Robust standard error in parentheses)		
*** p < 0.01, ** p < 0.05, * p < 0.1		

Simulation of prices, inflation and compensating policies

The scenarios concern the increases that Europe is currently experiencing in terms of inflation and gas prices- with the latter increase associated with the ongoing energy crisis (Steckel et al., 2022). For this exercise, the latest (2022) gas prices and inflation data for the 28 countries have been extracted from Eurostat. As a proxy to simulating inflation increase, the 2020 household incomes in each country were reduced with a rate equal to that of the inflation increase in that country. The three final scenarios additionally concern hypothetical compensating policies aiming to alleviate the effects of price and inflation increase. These compensating policy simulations are based on our model's social protection payments and energy efficiency parameters, because (as mentioned earlier) these represent measures that are already prominent in the energy poverty alleviation agendas. Each one of our scenarios is described below:

- Scenario 1 (S1): the price of gas is increased as in the beginning of 2022.
- Scenario 2 (S2): the price of gas and the inflation are increased as in the beginning of 2022.
- Scenario 3 (S3): the conditions of S2 remain true, while a compensating policy of increasing heating energy efficiency by an indicative rate of 10% is implemented.
- Scenario 4 (S4): the conditions of S2 remain true, while a compensating policy of increasing social protection payments is implemented. The increase rate is the same as the inflation rate.
- Scenario 5 (S5): the conditions of S2 remain true, while increases in both efficiency and social protection payments are implemented.

For S1 and S2, Figure A.1 demonstrates an average (across household types) increase in the probability of being energy poor. Afterwards, and by implementing the compensating polices, there is an average decrease in the probability. Therefore, increasing the social protection payments (S4) has a greater effect in decreasing the probability compared to increasing energy efficiency (S3). Moreover, when both efficiency and social protection payments are increased in S5, there is a join effect that decreases the probability the most. However, this is not adequate to fully alleviate the damaged occurred in S2.

Finally, from Figure A.1 it is evident that single parent and one person households are more energy vulnerable compared to other types, which is a plausible observation.

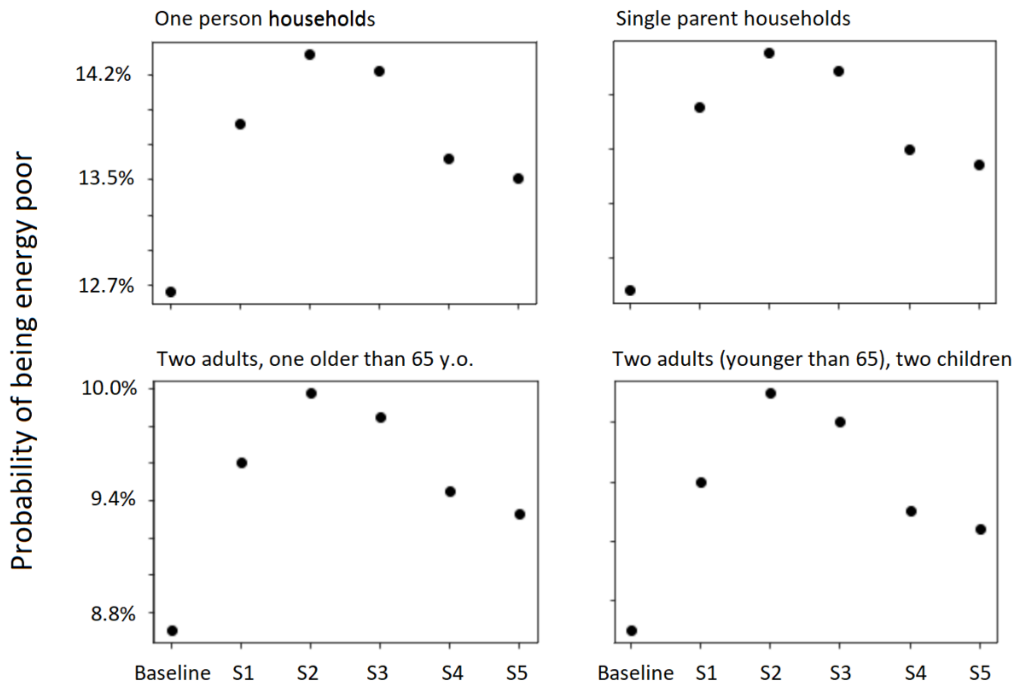

Figure A.1: Changes in probability of being energy poor in EU-28, for certain household types and under scenarios S1-S5.