Investigating the use of Machine Learning for Resource Intensive Agent-Based Simulation and Optimisation of Domestic Energy Retrofit Adoption

by James Hey

Thesis submitted to the University of Nottingham

for the degree of Doctor of Philosophy

supervised by

Dr. Peer-Olaf Siebers
Dr. Ender Ozcan
Dr. Paul Nathanail
Dr. Darren Robinson

July 5, 2022

Abstract

The existing domestic urban built environment contributes significantly to the environmental issues facing the international community. Whole house energy retrofits performed on this stock are a key tool for the mitigation of greenhouse gas emissions. As such, researchers have built a set of bottom-up retrofit adoption models to understand, model, and predict household energy use and retrofit decisions in this space. However, the existing methods are limited in their ability to estimate complex decisions for large housing stocks due to computational restraints of optimisation. While machine learning has been applied to the problem of retrofit optimisation for performance improvements, the methods still scale poorly when applied to bottom-up agent-based models due to the large number of heterogeneous problems to be solved. This thesis aims to advance this toolset by extending state of the art data science approaches to domestic retrofit decision modelling across urban housing stocks.

The major contribution focuses on a transparent method of obtaining rapid predictions for near-optimal energy retrofit solutions using deep neural network models. This process is referred to as surrogate optimisation due to the surrogate modelling techniques used to achieve it. The models are trained on a sample of near-optimal solutions generated using traditional surrogate energy models paired with optimisation techniques to obtain a data set of retrofit decisions for model training. This allows for rapid estimation of retrofit decisions based on both the physical characteristics of the dwelling and the social characteristics of the households that would not be computationally feasible using existing methods.

This process is initially limited to the single objective of net present value to model rational and self-interested agents. The process is then extended to a multi-objective problem by considering net carbon emissions savings. This allows manipulation of the objective function to expose the household emissions valuation, which represents the marginal financial value a household places on each ton of carbon. By training the surrogate optimiser with these values, it was possible to both generate Pareto Fronts and target a specific carbon value held by a household, a characteristic that is both measurable and simple to understand. An agent-based model was constructed using the survey data derived decision model with the household carbon value trained surrogate optimiser. This allowed the consideration of scenarios and policy at a scale and level of detail which, without this research, would have been unfeasible with the computational resources used. After demonstrating the surrogate optimisation technique we perform a novel analysis of Best-Worst Scaling survey data in an attempt to understand household decisions better, ultimately resulting in a retrofit decision trigger model based on the responses. The more realistic trigger model was used in conjunction with the surrogate optimiser, capturing more realistic investment decisions.

Acknowledgements

Firstly, I would like to thank my supervision team. I am deeply grateful to my primary supervisor Peer-Olaf Siebers, Assistant Professor at the University of Nottingham, who provided reliable and wise guidance at all stages of my research. Professor Darren Robinson who provided significant support and challenging critique during his time as LUCAS project leader and after the project was completed. Dr. Paul Nathanail for his incredible strategic perspective, helping me stay grounded on achievable and realistic aims while planning for long-term research goals. Finally, Professor Ender Ozcan who was always willing to find time to provide feedback and support.

I would like to thank those who offered up their precious time to proofread, critique, and improve my work. In particular, Katherine Rodríguez's pursuit of every supernumerary comma, confusing turn of phrase, or outright misuse of the English language, has undoubtedly protected my sanity in ways for which I cannot express enough gratitude.

I am grateful to my office mates from B36 for their academic stimulation and hours of challenging discussion, particularly Jonathon Thaler, Kwabena Amponsah, and Utkarsh Agrawal who kept me on my toes intellectually for this entire journey. Without the atmosphere of questioning and inquiry they created, conducting this research would have been a measurably drearier affair.

I would also like to acknowledge the support provided by my family. The advice, wisdom, and support of John Hey who was always just a phone call away on the many occasions that it was needed. His advice and support on day-to-day problems as well as long-term strategic planning gave me the boosts I needed at critical times to accomplish my research goals. The support provided by my Mother, Janette Medcalf was very different in nature but equally important, providing financial and morale boosts when they were needed most.

This research was funded by the Leverhulme Trust (RP 2013-SL-015) and supported by the University of Nottingham's Computer Science department.

Statement of Originality

I declare that the content of this work is my own and assistance received in preparing it has been appropriately acknowledged. This thesis has not been submitted for any other degree or other purposes. Some of the content of this thesis has been previously published within the publications listed in the text.

Contents

List of Figures

List of Tables

Table 1: List of Key Abbreviations

ABM	Agent-Based Model/Modelling
ANN/DNN	Artificial/Deep Neural Network
BPEC	Building Performance Evaluation and Certification
BWS	Best Worst Scaling
CDF	Cumulative Distribution Function
EPC	Energy Performance Certificate
EWI	External Wall Insulation
GA	Genetic Algorithm
GJ	gigajoule
HCV	Household Carbon Value
HSEM	Housing Stock Energy Model
HSRAM	Housing Stock Retrofit Adoption Model
IWI	Internal Wall Insulation
kWh	kilowatt hour
LCS	Lifecycle Carbon Savings
ML	Machine Learning
MAE	Mean Absolute Error
NPV	Net Present Value
RMSE	Root Mean Squere Error
ROI	Return On Investment
tCO2e	tonnes of Carbon Dioxide (CO2) equivalent
SAP	Standard Assessment Procedure
SEPM	Surrogate Energy Performance Model
SHAP	SHapley Additive exPlanations
SO	Surrogate Optimisation/Optimiser
TOID	TOpological IDentifier
WHRS	Whole House Retrofit Solution
WTP	Willingness to Pay

Chapter 1

Introduction

1.1 Motivation

There is significant scientific consensus that anthropogenic greenhouse gas emissions are leading to climate change [2, 3, 4]. The UK government has committed to a deadline of 2050 for net-zero emissions [5]. Domestic dwellings account for over a third of the national energy demand and approximately a quarter of total CO_2 emissions in the UK, with the majority of this energy demand being used for electric or gas based space heating [6, 7]. While significant progress has been made in the design of new buildings for efficient heat generation and retention, most of the existing stock will still be occupied by the UK government's deadline of 2050. Meeting this target will require strategies to transform the existing building stock with near-optimal combinations of different energy saving measures, known throughout this work as Whole House Retrofit Solutions (WHRS). In addition, to measure the effects of potential policies, as well as optimise any parameters they may introduce, high-quality retrofit adoption models are required.

While there exists a range of tools to analyse the impact and adoption of retrofit solutions in the existing stock, there are a number of issues with existing methods. Top-down building stock models are inflexible, based on aggregate trends, and are often dependent on rigid historical data, which leaves little room for detailed analysis of either input parameter changes or granular model outputs. Bottom-up retrofit adoption models benefit from significantly finer control over the simulation inputs, however, this increased level of detail comes with significant computational costs. This is particularly difficult when considering the heterogeneous nature of physical dwellings and the household decision makers who occupy them, as the selection of a WHRS would normally be performed for each dwelling by an expert assessor who would select a near-optimal WHRS to propose to the decision makers, based on their preferences. This process can be replaced with an optimisation procedure using both the physical and social properties of the household to model the decisions that could be made, but performing this detailed optimisation procedure for every dwelling across large building stocks is not computationally feasible.

The computational complexity introduced when using bottom-up modelling approaches can be considered in terms of a simulation possibility frontier. This is a visual representation of a conceptual problem: the trade-off between the level of detail within a simulation model and the technical feasibility of that model. As the quantity or complexity of phenomena captured within a given model grows, bringing it closer to the near-infinite complexity of reality, the technical feasibility of building and running the model diminishes. This is shown by the bounding curve in Figure 1.1 and captures the difficulty of bottom-up modelling approaches which involve complex operations, such as the large scale retrofit adoption modelling of a building stock.

Expansion of the simulation possibilities curve, as demonstrated by the dashed line, is to be expected over time, as Moore's law continues to drive the improvement of available computing capacity. However, the boundaries of feasible simulation can also be extended by improvements in modelling methodology which are able to capture a higher level of detail and/or diminish the computational cost of capturing given phenomena. The methodological contributions made in this thesis attempt to explore new simulation methods which are able to capture additional levels of detail while maintaining the technical feasibility of large scale models by integrating machine learning and novel agent-decision models into bottom-up agent-based models. These contributions will be made in the domain of WHRS-adoption modelling.

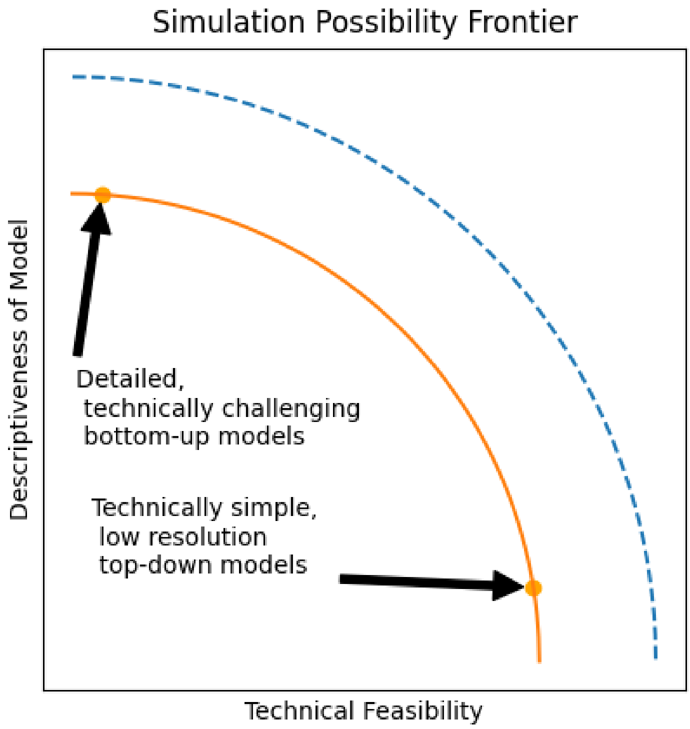

Figure 1.1: Simulation possibility frontier - the boundary of descriptiveness and technical feasibility combinations at a given moment in time.

1.2 Research Aim

The **research question** of this thesis is: How can Data Science techniques be used to expand the simulation possibility frontier, allowing for more descriptive bottom-up, large-scale energy retrofit adoption models while maintaining the technical feasibility of the simulation, in the context of an agent-based simulation of urban energy retrofit models?

The **research aim** is the statement formulation of this question: To expand the simulation possibility frontier using the Computer Science toolkit, to allow more descriptive bottom-up large-scale retrofit decision modelling without impacting the technical feasibility of the model.

1.3 Contribution Statements

In order to address the research aim and answer the research question stated above, the following contributions have been made throughout the course of the work presented in this thesis. It is put forward that the satisfactory demonstration of these contributions will answer the research question and demonstrate an original and significant contribution to knowledge.

Contribution 1 - To the best knowledge of the author, this thesis is the first to systematically investigate the integration of Surrogate Optimisation (SO) into an Agent-Based Model (ABM) to analyse energy retrofit adoption in urban housing stock. The integration of the SO technique into an ABM allows for increasingly rational, self-interested agents to be simulated at a scale that would otherwise be infeasible by allowing computationally cheap optimisations. The investigation in this thesis considers the implementation challenges, performance, and drawbacks of this method of ABM analysis.

Contribution 2 - To the best knowledge of the author, this thesis is the first to systematically investigate the extension of the principle of multi-objective Surrogate Optimisation for the analysis of domestic urban energy retrofit potential that includes a measure of households' Willingness to Pay for carbon mitigation. This has allowed for the conception of the ABMs created for Contribution 1 to relax the assumption of self-interest by considering environmentally conscious agents.

Contribution 3 - To the best knowledge of the author, this thesis is the first to combine a data-driven retrofit trigger model with a Surrogate Optimisation method. The decision trigger model takes survey data from participants to determine when retrofit adoptions are likely to be considered, as modelled in Contribution 2. This allows for models which conceptualise and include heterogeneous decision factors while maintaining intelligent and preference driven retrofit evaluations.

1.4 Publications

- Hey, J., Siebers, P., Ozcan, E., Nathanail, P., Robinson, D. "Surrogate Optimisation of Energy Retrofits in Domestic Building Stocks using Household Carbon Valuations" *Journal of Building Performance Simulation* (Under review)

- Hey, J., Siebers, P., Ozcan, E., Nathanail, P., Robinson, D. "Surrogate Optimisation of Housing Stock Retrofits using Deep Neural Networks" *Building Simulation and Optimization* 2020

- Siebers, P., Zhi, E.L., Figueredo, G., Hey, J. "An Innovative Approach to Multi-Method Integrated Assessment Modelling of Global Climate Change" *Journal of Artificial Societies and Social Simulation* 2020

1.5 Thesis Structure

Chapter 1 introduces the motivation and context of this thesis, lays out the broad research aim and the contributions presented in order to achieve the research aim.

Chapter 2 provides the context required to understand the work performed for the reader unfamiliar with the concepts involved. This will include discussions of agent-based modelling, machine learning, optimisation techniques and relevant areas of the built environment. While this will require reference to the literature on a broad range of topics, this will exclude areas most relevant to the specific contributions of this work, which will be addressed in Chapter 3.

Chapter 3 involves the presentation and discussion of the most relevant literature to highlight the context and contributions made to the fields involved. This includes a discussion of existing methods to evaluate the energy use and retrofit adoption of existing buildings and building stocks. The limitations and drawbacks of existing techniques are highlighted in Section 3.4, which identifies explicitly the research gap that this thesis attempts to fill.

Chapter 4 explains the approach taken to solving the research aim. This includes a broad description of the research methods which will be carried out in Chapters 5, 6 and 7 as well as a discussion and investigation of the primary data set to which these methods will be applied. Given the methodological contributions of this thesis, which are detailed more rigorously in later sections, only a top-down view of the methods will be discussed in this chapter. Specific implementation details, such as model selection and parameter tuning, are presented alongside the results of each method in their related chapter.

Chapter 5 lays out the creation of a single-objective surrogate optimiser to capture the behaviour of environmentally indifferent rational agents. These agents attempt to maximise the economic return from a retrofit investment. The Surrogate Energy Performance

Model (SEPM) and Genetic Algorithm (GA), both of which are used throughout the research, are also introduced and evaluated here as preliminary steps.

Chapter 6 expands the surrogate optimisation technique to capture the behaviour of environmentally conscious agents with a model of environmental contributions using multi-objective optimisation. Agents are instantiated with carbon valuations representing their willingness to pay per ton of carbon mitigated, which is used to form an objective function to perform optimisation. Carbon valuations are sampled to create a set of Pareto fronts describing near-optimal solution sets for each household, which are then used alongside the carbon valuations to train a set of preference aware surrogate optimiser models.

Chapter 7 investigates how to relax the stricter decision model assumptions made up to this point by introducing a pilot survey data. This includes estimating the significance of retrofit decision triggers among participants and then pair these to the attributes of the synthetic population, capturing the retrofit decision process in a higher level of detail. This trigger model is combined with the Surrogate Optimiser introduced in Chapter 6 to create a detailed and computationally inexpensive retrofit adoption ABM.

Chapter 8 consists of two sections to complete the thesis. The first section summarises the work conducted, highlighting the contributions, and discusses their significance. The limitations and points for future work are also laid out here. The final section contains concluding remarks, giving some overview of the purpose, aims and achievements of the presented work.

Chapter 2

Background

This chapter will serve to provide the background information required to understand and contextualise the contributions found in later chapters. Given the interdisciplinary nature of the research, some areas in which contributions are not made are still highly relevant and covered within this section. This chapter will also explain and formalise some common language and practices used when building and reporting different types of models present throughout the work, especially where different disciplines use identical or similar language for different concepts. This chapter is supplemented by a discussion of the most relevant literature found in Chapter 3, which will identify the gaps in the existing literature and highlight the areas of literary contribution.

2.1 Modelling and Simulation

2.1.1 What is Modelling?

We will define a model as a man-made construction attempting to imitate aspects of a target system. This definition is quite broad, allowing it to encompass all types of model to be considered in this thesis, and will allow us to identify: what the target system is, the aspects of it that we will be imitating, the degree to which we will imitate them, and which aspects we are not imitating, whenever a new model is introduced. It is possible, and I believe important, to do this for any well-understood model. Consider the example of a mannequin designed for displaying clothes in a shop window. Here, the target system is of the clothes worn by the potential buyer, imitated by the clothes on the plastic model. This is simulated by the approximate size and shape of the model being of standard human proportions, making clothes rest on the model in a similar way as they may on a person. However, many aspects of the target system are unmet. The size and shape of the model, while representative, are very unlikely to match the exact characteristics of the buyer. The model does not move. It lacks hands and feet. The texture of the skin is smooth and synthetic. In these ways and countless others, the model deviates from the target system. Some deviations, such as its inability to act and think for itself, are not relevant to the purpose of the model. Other deviations, like the representatives of the model to a given shopper, do impact the effectiveness of the model, but in a way that is either impractical or impossible to imitate. There are also some deviations, such as the rigidity of the mannequin which deviate from the target system but are desirable properties of the model, as this rigidity allows for any clothes applied to remain in place for long periods of time. It is also worth noting that in this example the model is the mannequin itself, and does not include the clothes that it may currently be wearing, as this would be an instance of using the model to perform a simulation.

2.1.2 What is Simulation?

We will define simulation to be the use of a model to understand or predict the behaviour of the target system. This is done by observing the model in its imitations of the target system. We can consider simulation to therefore be the active, in contrast to the model which is the passive. Taking our mannequin as the model, for example, our simulation would be to dress the model in clothes and observe how they appear. Here we can see one of the aspects of simulation: parameter variation. We are able to experiment with our model by simulating a variety of different clothing combinations. When we discuss simulation, it is therefore important to discuss these parameters: the scenarios we place on our model which influence the observed outcome. Some, like the choice of clothes, may be the primary factors of interest that we vary with experimentation. Others, like the lighting in the room or the pose of the mannequin, which do influence our observations, may need to be held constant with their choices justified. We must therefore ensure, when discussing a specific simulation that we understand both the explicit experimental factors and the less obvious environmental assumptions put in place during that experiment.

2.1.3 Modelling and Simulation Paradigms

Given our extremely broad definition of modelling, it is important to be more specific in describing the subsets of models that we will be focusing on. There are many ways to build models of a target system, and there is little benefit in considering them all. Instead, we will first discuss the difference between top-down and bottom-up approaches to modelling complex systems. We will then look at the use of agents to build bottom-up models, followed by a discussion of statistical models. Statistical models will be examined in more depth in Section 2.2, as they are a core component of machine learning techniques.

2.1.3.a Top-down vs Bottom-up Approaches

Top-down approaches attempt to capture the entire system's behaviours broadly using aggregated values. A top-down approach to modelling heating energy use in a city may, for example, use an aggregate of the historical demand. A growth rate could be applied, based on the expected rate of building new dwellings, expected changes in technology, decarbonisation of the grid, etc. The individual buildings themselves would not be considered, and may not be required for the simulation desired. The limitations of top-down modelling come from the level of detail that certain phenomena can be captured in, as there is significant abstraction and loss of detail when using aggregates.

Bottom-up approaches, in contrast, attempt to capture the behaviour of the target system by breaking it up into smaller sub-components to be modelled separately. The behaviour of these components are then aggregated to measure the behaviour of the entire system. This can allow for the inclusion of behaviours or interactions that are not modelled by the top-down approach. Let us consider a bottom-up approach to modelling heating energy demand in a city. The modeller can individually model the demand used in each building based on the individual characteristics of that dwelling, the number of people in the dwelling, and other heterogeneous details that are known to influence energy demand. When aggregated for all dwellings, this should give the demand for the building stock as a whole. The challenges of bottom-up approaches begins to become apparent here, as a significant amount of data about the sub-components that make up the model may be missing. While there are methods to overcome these, they often require the introduction of error through the assumptions made.

The main advantage of bottom-up approaches is the additional level of detail that can be captured by using disaggregated analysis. These approaches expose those low-level phenomena to the modeller as a source of experimentation and analysis. In our hypothetical energy use model, for example, a bottom-up approach may allow us to identify which buildings or types of buildings are more likely to contribute to energy demand, allowing policy makers to design more targeted measures based on details that are obscured in top-down approaches.

A key disadvantage of bottom-up approaches is the computational complexity in running them. In the best case, where sub-components do not interact significantly, they scale linearly with the number of sub-components. Even in this case, the number of heterogeneous components may be large or the modelling complex [8]. If these components interact, then the complexity will increase exponentially, rendering larger systems infeasible to model as they scale.

Given these approaches can be used to model the same target systems we would expect them to yield comparable results, although in practice there has been a gap [9]. This suggests errors in one or both of the approaches when building models. Which approach is more error prone is context dependent. Top-down models may introduce errors by excluding low-level phenomena that are captured in bottom-up approaches. In contrast, bottom-up models introduce more detailed phenomena and often many more parameters, each of which could be prone to its own error. We used a method of reconciling the two approaches when developing a hybrid integrated climate assessment model. We added a more bottom-up approach to an existing top-down model, but validated the outputs against the top-down model [10]. Once validated, the additional parameters exposed by the bottom-up approach can be adjusted to achieve the more detailed analysis that bottom-up approaches grant.

Top-down and bottom-up approaches are not necessarily mutually exclusive or rigidly defined paradigms. Models can be considered as existing on a spectrum with no model being purely bottom-up or top-down in nature, but rather capturing more or less detail about the sub-components included. Sometimes these models are referred to as hybrid models. Consider our example of a bottom-up energy demand model: the households are modelled heterogeneously, but we may be using a top-down approach when considering certain behaviours. We are likely to model the households' behaviour rather than each individual member. Due to the near infinite complexity of the target system, at some stage modellers are obliged to draw a line and begin using aggregated or simplified behavioural mechanics. Hybrid modelling techniques may also include procedures such as sampling, in which sub-components of a larger system are sampled, modelled in a detailed, bottom-up manner, then extrapolated, in an attempt to understand the entire systems. Housing stock models of this type are discussed in Section 3.2.

2.1.4 Agent-Based Models

Agent-Based Models (ABM) are a bottom-up approach to modelling, performed by capturing the simplest fundamental behaviours and encoding them into agents as part of a system. The advantages of ABMs are three-fold: they allow for emergent phenomena; they are a natural way to describe a system; they are flexible [11]. Emergent phenomena in ABMs occur when a relatively simple set of rules between agents results in more complex observable behaviour. The phenomena may be more complex than can be captured in a top-down manner, allowing observations that are otherwise difficult to measure or model. ABMs are useful for geographic models where the location of agents in physical space is relevant [12]. There is sometimes confusion about the difference between Multi-Agent Systems (MAS) and ABMs [13]. This confusion is perhaps unsurprising, as they both implement multiple interacting agents which fit the classifications laid out above. However, for the sake of consistency in this thesis, we will consider ABMs those which rely on a large number of (relatively simple) interacting agents in a system designed to represent the real world. In contrast, MAS rely on a smaller number of agents, often in a more abstract domain.

Before going any further, it is important to determine the properties of an agent, so as to identify when and when not, we are considering an ABM. Russell and Norvig's definition of agents is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through effectors [14, p.31]. This broad definition encompasses humans, software agents, simulation agents, and robotic agents whilst excluding anything we would consider evidently not an agent. This makes it a beneficial working definition. Given that we may observe seemingly disparate types of agents falling under this definition, there is utility in considering some classifications of agents so that we are mindful when designing or discussing that which we believe to be an agent.

Russell and Norvig [14] begin their agent classification process by identifying four agent types based on the nature of their environmental interactions, although additional behavioural profiles have since appeared in the literature. The first agent type introduced is reflex agents. Reflex agents are designed with a direct mapping between their observed environmental state and the actions they are to take. The second agent type, reflex agents with state, uses not just the observation of the environment but also an internal recording of state to determine state-action mappings. The final two agent classes are goal-based agents and utility-based agents. Goal-based agents are given a desired state for the world, instead of the direct mapping between states and actions of previous types, and take action in an attempt to close the gap between the current world state and the desired one. Utility-based agents are something of an extension to these, with the designer assigning utility to positive states of the world and disutility for negative ones. This allows for conflicting objectives to be balanced and for the issue of partial rewards for incomplete goals. As with any such classification system, there are potential flaws that can be identified in these types. One could argue, for example, that goal and utility agents must, in order to ultimately act, have a direct mapping between state and action and therefore are simply more complex versions of reflex agents. It could also be argued that reflex agents have a goal that has been predetermined by the designer. Nonetheless, these classifications can be useful for agent design and even for determining if the sub-components of a system are indeed agents.

2.1.5 Microsimulation

The definition of agents laid out above is quite broad, with any entity that observes and acts on its environment being considered to have agency. There have been classification of modelling methods which draw a distinction between agent-based modelling and microsimulation. Microsimulation models are associated with one-directional impact of policy interactions, while agent-based models focus on emergent behaviour from more complex inter-agent relationships [15]. The line between agent-based models and microsimulation is also not always fixed, with some examples of hybrid models that contain features distinctive to both definitions [16]. Given the objectives laid out in this thesis, the distinction between agent-based and microsimiulation models is not of primary significance, as the focus is on increasing the scale and intelligence of decisions made by agents, findings which will hold regardless of whether the agents are interacting with other agents or acting independently in response to policymaker decisions.

2.1.6 Mathematical Models

$$q = -k \cdot \frac{T_2 - T_1}{L} \tag{2.1}$$

Equation 2.1: Mathematical model of a thermal flux (q) through a surface of thickness L, a thermal conductivity of k, and a temperature differential of T_2-T_1.

A mathematical model attempts to capture the behaviour of the target system in terms of closed-form mathematical expressions. These models can be constructed from observation or derived from theoretical models of the working of a target system. An example of such a model would be a model of thermal conductivity shown in Equation 2.1. This mathematical model can be used to calculate the rate of conductive heat flow through a surface, one of the many elements that make up building energy simulations discussed in Section 2.3.4. Indeed, it is common for larger simulation models to be constructed either entirely or partially of closed-form mathematical models which are then run together to perform simulations. This is most evident in system dynamics models, which are defined in terms of stocks of attributes and the flows between them, which are expressed as rates of change. Using feedback loops to link these mathematical components, complex non-linear stateful systems can be simulated over time [17].

Analytical models are a form of mathematical model constructed from observations of data. Generally, the structure of the model is chosen by the modeller, based on observations of data or an understanding of the target system. The data observations are then used to fit the models using a form of optimisation algorithm to find the most appropriate parameters. These analytical models are the basis for supervised machine learning, discussed in Section 2.2. This can be most clearly seen in linear regression models, discussed in Section 2.2.2.a, as the closed-form expression is explicitly defined in a linear form, compared with models such as neural networks which are less transparent.

2.2 Machine Learning

Machine Learning (ML) has been used as a somewhat nebulous term for a broad range of techniques used in computer and data science. The field relies heavily on statistics, as many techniques rely on the automated construction and use of statistical models using data [18]. Introductory texts generally contrast machine learning techniques against traditional algorithmic problem solving used in computing [19]. While algorithms are used to generate the models used in ML applications, they are generic and can be applied to disparate problems by providing different data on which to train them. The algorithms used in machine learning are often further divided into categories depending on the nature of the data or processing which occurs:

- Supervised learning relies on a data set with known labels which can be used to fit a model. Examples include classification algorithms and regression models.
- Unsupervised learning attempts to detect patterns in unlabeled data. Examples include clustering algorithms and automated anomaly detection.
- Semi-Supervised learning relies on using a combination of labelled and unlabelled data to form a better model than just using one or the other alone.

- Reinforcement learning attempts to train an agent to devise a strategy to maximise a reward function within an environment.

The first class of machine learning techniques, supervised learning, requires a training data set in which the input and output data is known. This is a particularly robust method of training as, unlike in unsupervised methods, a test set and validation set can be retained before training, giving an indication of the performance of the model. These techniques are popular for classification and regression problems using methods which will be discussed in Sections 2.2.1 and 2.2.2. When talking about supervised learning models, it is worth breaking down the data into the input features and the target features. The input features, sometimes referred to as independent variables, are provided to the model at both training, testing, and deployment stages and contain the information used by the model to make predictions. The target feature(s), sometimes known as dependent variables, are the desired predictions of the model. At the training, validation, and test stages, these are known to the system and referred to as labels. When the model is deployed, only the input data is known and the model is used to predict the target feature(s). When discussing supervised learning models in this thesis we will identify the input and target features explicitly.

The second class of machine learning techniques, unsupervised learning, attempt to extract meaning from data without the use of known output labels. One common problem set involves clustering, which groups relevant data points together based on some proximity measurement learned by the algorithms. The clusters can either be analysed by the modeller to extract meaning, or used as input for a further model or decision algorithm. Both of these uses of clustering are explored in Chapter 7, and a discussion of clustering algorithms are explored in Section 2.2.4.

The third class of machine learning methods are semi-supervised techniques. These are hybrid methods that combine a labelled training set with an unlabelled data set in an attempt to achieve better performance than either of the techniques used alone. These algorithms are often adapted from supervised or unsupervised methods, using the additional data to supplement the algorithm in some way. These techniques can be very useful when obtaining labelled data is difficult or impossible, however, they will not be considered much in this thesis. This is because additional training data can be generated through simulation where needed. A very thorough discussion of the theory and application of semi-supervised learning can be found in [20].

The final class of ML techniques, reinforcement learning, would appear to be a natural match for an ABM as they involve learning through interactions with a dynamic environment. Indeed, this is a popular technique that has been used to grant more intelligent behaviour to utility and goal-based agents [21]. It is because of the popularity of these methods that we have largely excluded them from the scope of this thesis, as they are already well adapted to adding certain behaviour types into ABMs.

2.2.1 Classification Models

Classification models are a class of supervised machine learning techniques designed to identify label(s) associated with given input data. An example of a classification problem is the automation of spam email detection in which the input features would be the contents of the email (alongside any metadata available), and the target feature would be a boolean label identifying the email as spam or not spam. This type of model will be used when determining the insulation materials suitable for a given application in Chapters 5 and 6, as these materials take on discrete classes. There are a variety of methods for performing classification. We will consider a range of the most relevant methods to the research below.

Decision trees create a rule set to break the training data down into smaller and smaller subsets until each point can be classified based on the rules. This method is particularly powerful in capturing nonlinear relationships: bifurcated data sets or data with more outliers, that could impact other models more severely. Decision trees also benefit from being human readable, providing some insight into the most discriminatory rules. Random forests are an extension of decision trees in which a collection of different trees are initially trained. In deployment, the modal class from the forest is used, making this a form of ensemble classifier. This attempts to solve the overfitting problem that can arise from decision trees.

Binary logistic regression (also known as logit) is a classification method adapted from statistical regression models of the type discussed in Section 2.2.2.a. While regression models are normally more suitable for continuous target features, logit regression is adapted to using binary labels for training. This type of model is useful when wishing to extract insight as the coefficients are proportional to the significance of the independent variable. We use this as both an exploratory and predictive modelling technique in Chapter 7 when analysing the contributing factors of binary decisions made by survey respondents. It is also possible to use standard regression techniques for classification if classes can be determined by the binning of continuous variables. A regression model to predict household income, for example, could naturally be used to classify low income households with a given income threshold.

Support vector machine classification uses a kernel function to transform data into a higher dimension, then bisects it using hyper-planes which minimise the separation margin between classes. The same transformations are performed on novel data, with the trained hyper-planes used to determine the class of the new data. This technique performs well with high-dimensional data and it is suited to situations where small sample sizes are necessary, although can struggle with noisy data where classes are less clearly defined [22].

The K Nearest Neighbors (KNN) algorithm is a simple but effective algorithm that can be used for both regression and classification tasks. Being non-parametric, it does not rely on knowledge of an underlying distribution and is, therefore, less likely to train bias should assumptions of the distribution be incorrect [23]. When predicting with KNN, an unclassified input is compared with k's (a parameter of the algorithm) closest values in the training set and assigned the class of the plurality of those neighbour's classes. The definition of closeness is important for this algorithm and may require careful pre-processing to ensure that the values are normalised and that irrelevant or highly noisy features are removed to prevent distortion. Dimensionality reduction techniques, such as principal component analysis, are helpful with dealing with some of these problems [24].

2.2.2 Regression Models

In contrast to classification, in which discrete labels are applied to data, the target feature of regression models is continuous. This makes them suitable for modelling many real-world systems, such as predicting income. Where models require continuous outputs, a regression model will be used in this work. This includes the SEPM implemented in Chapter 5, which is used to predict the energy demand of a property with a given set of dwelling attributes.

2.2.2.a Linear Regression

$$y_i = \beta_0 + \beta_1 x_{i1} + \beta_2 x_{i2} + ... + \beta_p x_{ip} + \epsilon_i \tag{2.2}$$

Equation 2.2: Multiple linear regression with p explanatory variables.

Linear regression is an attempt to fit data into an analytical form demonstrated in Equation 2.2, where the target feature of a given observation (y_i) is modelled as a linear combination of the products of p's explanatory features (x_{ip}) and a coefficient ($\beta_i p$), as well as an error term (ϵ). The error term represents the distance between the model's estimated value for the target feature and the true value. The process of training a linear regression model is to optimise the coefficient variables (β_p) in order to minimise a loss function, typically the sum of the squared of the error terms in the training set.

This model formulation has some key advantages, not just in the relatively cheap computational cost of construction, but also in the transparency of the trained model. Once trained, the coefficients make it clear how much each explanatory variable contributes to the model and they are often used to provide insight into the significance of these variables on the target system.

There are some issues with linear regression formulations, mostly based on the assumptions required for the model to fit well with this method. If the target system is not linear then the model is likely to fit poorly, with coefficients giving little insight into its workings and high error. There exist variants of linear regression models that attempt to handle non-linearity by adding polynomial terms that still combine linearly to form a model. Linear regression models also assume homoscedasticity, meaning that the variance of the errors remains constant across the entire domain, a property that must be tested for an accounted for if present. It is also important that variables used in linear regression are not correlated with each other, either directly or through a missing latent variable. These forms of under-specified model result in biased models spurious coefficients, and therefore linear regression model's require careful specification to ensure they are reliable.

2.2.3 Artificial Neural Networks

Schmidhuber (2015) [25] reviews the history of Artificial Neural Networks (ANN), noting that the concept dates back to as early as the 1800s. The technique uses artificial neurons designed to trigger under certain conditions, which can be used to map inputs to outputs by linking neurons into a network. The neurons are arranged in layers, with input neurons activated in correspondence to the input data (which may be training data, test data, or data from the target system). The input neurons are then connected to a hidden layer of the network via weighted connections. This hidden layer of neurons then connects to the final output layer with another set of weighted connections. These final layers can be adapted to perform either classification or regression by altering the activation function of the final layer.

Neural networks are trained by adjusting the weights of connectors between layers. The standard technique for network training relies on a set of processes called backpropagation and gradient descent. The network is initialised with untrained weights and the training data is run through the network. The network outputs are compared to the labels and the amount of error (defined by a loss function). Gradient descent is the method of adjusting the weights of the connections in the network to reduce the size of the error. Backpropagation is the process of working backward through the network to calculate the gradient by which weights within the network should be adjusted. This gradient descent through backpropagation is performed iteratively, in an attempt to reduce the loss function of the network. It is possible to overfit neural networks, reducing their predictive capability when presented with unseen data. To detect overfitting, a validation stage can be included, which uses separate validation data during training to measure the errors on unseen data. This allows the modeller to end training when the performance of the model begins to suffer from overfitting.

A Deep Neural Network (DNN) extends the ANN principle through the inclusion of multiple hidden layers, creating greater separation between the input and output neurons, resulting in a larger set of weightings and more paths through the network. The theoretical purpose of these additional neurons is the ability to break the input features into more complex, discovered features that can then be acted upon. This includes positive performance when data comes from a complex or heterogeneous distribution, in contrast to the weakness of models like linear regression, as discussed in Section 2.2.2.a. The additional layers allow, for example, a different route to be taken through the network for differing building types, which can result in more intelligent mappings from different input types.

Adding depth to ANNs was considered in the 1960s and 1970s but did not become practical for use until after the development of backpropagation as well as hardware improvements [25]. By the 2000s ANNs of various types began to outperform other methods in domains such as image classification, speech recognition, and many other applications. As we will discuss in Chapter 3, this method has recently gained traction for surrogate modelling in the built environment literature.

DNNs have a selection of hyperparameters chosen by the modeller in order to build well-performing models. Where hyperparameters are selected through this research, we will identify the choice and, where necessary, provide justification for that particular parameter. This can be done by the selection of sensible defaults that work well for similar problems or through a hyperparameter tuning process that includes training multiple models and comparing the results.

One important hyperparameter when training a DNN is the structure of the network itself. While the input and output layers are likely determined by the shape of the data being handled, the number and size of the hidden layers must be selected. The optimal configuration is highly dependent on the problem domain, with problems like image recognition benefiting from a large number of layers on neurons, often with complex structures like convolutional layers. Lower dimension classification problems, on the other hand, generally require fewer total neurons. While some intuition can be used to determine the optimal network structure, based on an understanding of the underlying features or dimensions of the problem, it is often necessary to discover a well-functioning structure through a tuning process. When it is necessary to do so in this thesis, a starting point will be taken from similar problems in the literature, and then a range of structures will be tested and their performance compared.

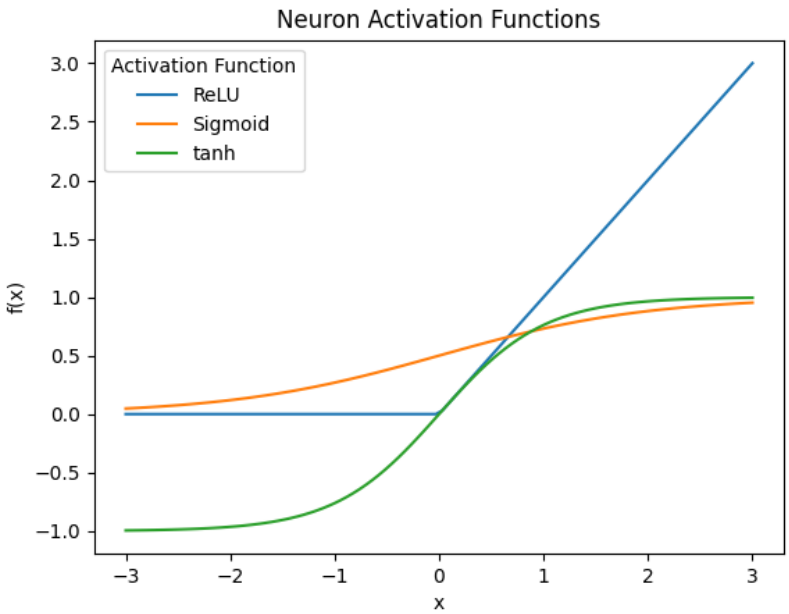

Figure 2.1: Comparison of Rectified Linear Unit (ReLU), hyperbolic tangent (tanh), and Sigmoid activation functions.

The activation function defines the non-linear firing of a given neuron based on the inputs. There are various properties of activation functions that have effects on the network. Non-linearity is important as networks with linear activation functions are equivalent to single-layer networks, mitigating many of the benefits of this choice. Popular activation functions include the Rectified Linear Unit (ReLU), the Sigmoid function, and the hyperbolic tangent (tanh) function. These functions, as shown in Figure 2.1, are both non-linear and differentiable, making them suitable for DNNs trained with backpropogation. The effectiveness of a given activation function will depend on the problem and may need to be tested. While the Sigmoid and tanh functions were traditionally more popular, ReLU has been found to perform well in tasks with a large number of neurons such as signal processing [26]. As neurons with ReLU activation can be fully deactivated for all $x < 0$, they are less computationally expensive to train, allowing for larger networks with more sophisticated pathways.

2.2.4 Clustering Techniques

Clustering techniques are unsupervised methods of grouping data points based on some metric of similarity. This can be done with the intention of learning something of value through cluster analysis or used to give structure to data for the purpose of further processing. Clustering has many parallels to classification (discussed in Section 2.2.1), with the primary difference that while classes are initially provided as labels in training data, clusters are not defined or provided but rather generated as part of the process.

The unsupervised nature of clustering, where there are no training labels to evaluate model performance, means that different metrics must be developed both to guide cluster formation and to analyse the quality of a terminated clustering process. Well-formed clusters yield two key properties. First, points within a cluster are similar (compact clusters). Secondly, points from different clusters are dissimilar (distinct clusters). Measures of these two fundamental properties can help evaluate the cluster quality and help validate the number of clusters. If, for example, the clusters discovered are compact but not highly distinct, this may be an indication that there are too many clusters.

There are several methods of creating and assigning clusters. The k-means clustering algorithm, initially laid out by J. MacQueen [27], takes the expected number of clusters (k) as an input parameter. Centroids are selected by an initialisation procedure, either by randomly assigning all points to a cluster or by selecting existing points at random to act as initial centroids. All points are then assigned to the closest centroid. The algorithm then iterates between recalculating centroids and reassigning points to the closest centroid, until the assignments no longer vary between iterations.

Fuzzy clustering methods differ from standard methods in one key way. Instead of attempting to assign a cluster to each point, these algorithms apply a degree of membership to each data point. The most common method for fuzzy clustering is the c-means algorithm as an adaptation of k-means clustering, which relaxes the constraint that cluster membership be binary and mutually exclusive [28, 29]. Interpretation of these fuzzy values is a common point of discussion. Fuzzy memberships of a given data point are not required to sum to one, making a probabilistic interpretation flawed. Instead, fuzzy membership can be considered the degree of membership a given point has to a given cluster, with 1 representing perfect point membership and 0 representing complete dissimilarity from a given cluster.

2.3 Optimisation

Optimisation is the search for the minimum (or maximum) value of a given function, as well as the inputs required to achieve that value. A generic form of this can be found formalised in Equation 2.3 which aims to minimise f with respect to x and y while satisfying a constraint, which in this case is a function of both input variables. Given the nature of this work, we will refer to the input variables of optimisation as decision variables, so as to differentiate them from the other input variables of an optimisation process, which are not always decision variables[1]. When performing optimisation problems in this thesis, they will be formalised like this to the extent possible; although complex optimisations may have black-box functions which cannot be laid out fully, or long, and sometimes unclear, sets of constraint functions that require shortening or simplifying.

In the example, we have decision variables x and y. The combinations of x and y can be referred to as the search space, solution space, or input space. When dealing with a constrained problem, we may also discuss the feasible solution/input space, which is the combinations of x and y which satisfy the constraint $g(x, y) \geq \epsilon$. When dealing with real values, the search space may be of infinite size. For combinatorial problems, the search space size is likely to be calculable and, when performing optimisations in this research, this will be done where relevant.

In this example, $f(x, y)$ is a generic function of two arbitrary decision variables, usually referred to as the objective (or cost) function. In reality, the system attempting to be optimised may be very complex, with a large number of decision variables, multiple objectives, and possibly uncertainty in outcome. In this thesis, the objective function value of a candidate solution is likely to be measured through a simulation run. Optimising based on simulations requires consideration of a few additional factors. If the simulation is stochastic, then a single evaluation is insufficient to determine the validity of a given solution. This can be resolved through repetition and statistical tests to state confidently

[1]Consider a building energy simulation, for example. The insulation thickness of a building may be a decision variable but the simulation has many other input variables over which we are not optimising. These non-optimised variables shall be referred to as input variables, but not decision variables.

that one solution outperforms another. Simulations are also likely to be time-consuming, often making the application of any exhaustive optimisation approaches infeasible for large search spaces. The simulations dealt with in this work are of this type, taking a long time, even in a single run, for the evaluation of a candidate solution, and they cannot be modelled or approximated analytically either. As a result, optimisation techniques falling into the (meta)heuristic category are preferred, to obtain a near-optimal solution within a reasonable amount of time as discussed in Section 2.3.1.

$$\begin{aligned} \min \quad & f(x,y) \\ \text{s.t.} \quad & g(x,y) \geq \epsilon \\ & x \in \mathbb{R} \\ & y \in \mathbb{R} \end{aligned} \tag{2.3}$$

Equation 2.3: Generic optimisation problem with two input variables and a single constraint.

2.3.1 Heuristic Optimisation Techniques

While some analytical functions can be solved to obtain a proven optimal value or searched exhaustively to determine a global optimum, there are many instances of complex systems in which this is not possible. These problems can be of enormous complexity when the search space (number and variability of inputs/decision variables) is very large and the objective function relating inputs to outputs is complex, disorderly, or computationally expensive. When an exhaustive search is deemed infeasible, optimisation heuristics can be used. Heuristics are a rule of thumb method that aims at finding a local optimal or near-optimal solution. Heuristics are problem dependent and must be tailored to the problem domain at hand.

Metaheuristics are also used in optimisation, which are general-purpose high-level optimisation/search algorithms that also require tailoring to the given domain [30]. Sörensen and Glover [30] classify metaheuristics into three groups: population-based, local search, and constructive approaches. Examples of local search metaheuristics, which construct solutions and then perform alterations upon them, include tabu search and simulated annealing. Constructive approaches, such as large neighbourhood search, create solutions from partial solutions, rather than altering complete solutions as would be done in local variations. A Genetic Algorithm is a very well-known example of a population-based metaheuristic, keeping a number of potential solutions under consideration at any one time. While the high-level guidelines that a metaheuristic optimisation employs are not domain specific, most have their own set of parameters that require tuning for best performance in a given search space. These metaheuristic parameters are referred to as *hyperparameters* to differentiate them from the parameters of the domain setting explored.

In broad terms, metaheuristics attempt to balance the exploitation of promising solutions with the exploration of a broader search space. Highly exploitative techniques are often faster to find local optima, but may miss better performing local optima due to a narrower search space [31]. In contrast, more exploitative techniques inspect a wider range of the search space but may take longer to converge to a near-optimal solution, if they are able to do so at all. These conflicting principles are sometimes referred to as diversification vs intensification, exploration vs exploitation, or breadth vs depth. The process of designing an effective metaheuristic includes designing algorithmic components with appropriate parameter settings, balancing these conflicting principles of exploration and exploitation during the search process.

Due to the heuristic nature of these algorithms, it is not guaranteed that globally optimal solutions will be found in all cases. The same algorithm may find different solutions each time it runs, due to stochastic elements. As such, we refer to solutions found from heuristic optimisation as 'near-optimal' as they provide no guarantee of global optimality. This also requires care when comparing an algorithms' performance, or even different parameter settings on the same technique. This can be overcome with repetition of the algorithm and comparison of the distribution of solutions generated.

2.3.1.a Hill Climbing with Local Search

Hill climbing is a greedy algorithm that performs a local search around a point in the solution space, adopting those which result in the best performance before continuing the local search. This method depends upon a domain-relevant heuristic of locality to allow for the exploration of neighbourhood solutions. Hill climbing can quickly resolve to a local optimum in a given region of the search space but can suffer from being too greedy by failing to take minor downhill steps that may result in converging on a better local optimum, located in a different region. There exist some variations that attempt to resolve this issue, such as repeat hill climbing which attempts to initialise from a broader range of locations across the search space to improve the chances of locating a near-optimal solution; or stochastic hill climbing, which adds random moves to the uphill steps, allowing for some exploration of the search space.

2.3.1.b Simulated Annealing

Simulated Annealing is a metaheuristic designed to imitate a process from metallurgy that involves the controlled raising and lowering of temperature [32]. It is similar to hill climbing but using a simulated system temperature to control which random local searches take place, with solutions being accepted provided they perform better than the current system temperature. During hot periods, this allows for exploration of a broader set of solutions but as the system cools, according to a cooling schedule, the search becomes greedier in an attempt to exploit the breadth achieved during the hot period. The cooling schedule usually contains a set of model hyperparameters which state of the art techniques automate in an attempt to increase the speed of convergence and quality of solutions [33].

2.3.1.c Genetic Algorithms

Genetic Algorithms (GAs) are a metaheuristic inspired by natural selection and genetics [34]. The process involves codifying the search space of solutions into a genome/chromosome, traditionally a string of binary digits - although this can be done with a variety of techniques. A random population of candidate solutions, consisting of chromosomes, are initialised and their performance in the search space tested based on the given objective function, referred to as fitness function. In each generation, the best performing individuals with the best fitness values are favored for reproduction and mutation. Depending on the implementation, mutation may include splicing successful genes together, applying minor mutations that make small perturbations, or major mutations that make large perturbations in the chromosomes. This way new individuals are produced, then a replacement scheme is applied to decide which individuals survive to the next generation. Poorly performing genes are expected to disappear from the specimen pool each evolutionary cycle through the replacement. This evolutionary cycle continues until the termination criteria are satisfied. A generic overview of a GA can be found in Section 2.2. Genetic algorithms show improvement over generations but, by applying mutations and keeping a variety of solutions in the pool, they attempt to avoid the premature convergence that occurs with greedy methods.

The core stages and components of a GA are as follows:

1. Create an Initial Population: An initial set of candidate solutions are created using a method, often randomly selected.

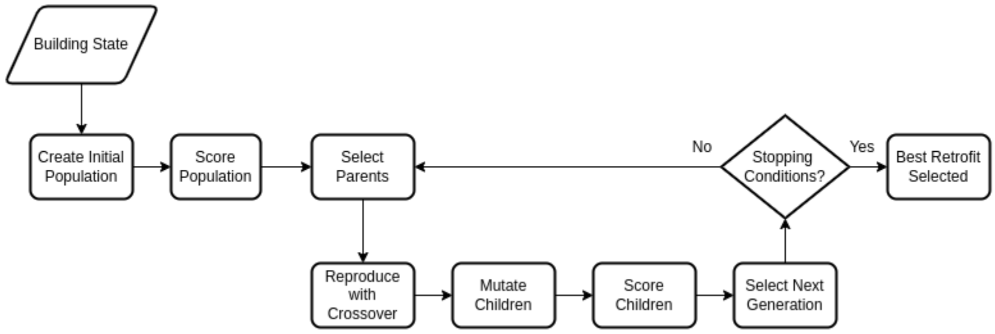

Figure 2.2: Process flow diagram of a generic GA.

2. Fitness Evaluation: This involves the scoring of candidate solutions in the population, based on the objective/fitness function.

3. Mate Selection: This is a method of selecting which candidate solutions will act as mates in producing the next generation of candidate solutions. Mate selection usually takes the fitness of the candidate solutions in a given generation into account, in an attempt to create pressure towards optima. It may still have a mechanism for maintaining solution diversity.

4. Reproduction with Crossover: This involves a strategy of generating new solutions based on the performance of existing solutions. Crossover involves a method of combining multiple (usually 2) solutions/individuals to form a new solution.

5. Mutation: This stage involves making minor local changes to the genes of a newly created solution(s)/individual(s). This adds a degree of local search around solutions.

6. Replacement: This involves selecting which members of the prior/new generation will continue onto the next. Replacement strategies can balance exploitation (selecting high scoring solutions) and exploration (leaving some poorer performing solutions).

Where a GA is used in this thesis, the strategies used will be discussed and justified when relevant. These optimisation parameters are referred to as hyperparameters or settings. Specific strategies may come with their own sets of hyperparameters, such as the tournament size when using a tournament mate selection strategy. Given the large number of hyperparameters and the stochastic nature of heuristic techniques (which require repetition and analysis of alternate combination settings), hyperparameter tuning can itself become computationally expensive.

2.3.2 Optimisation with Multiple Objectives

While optimising a single objective with respect to a set of decision variables is sometimes desirable, many real world problems require several different and competing objectives to be handled simultaneously, formalised generically in Equation 2.4 which shows a problem with two objective functions. Consider for example a building energy retrofit optimisation in which a homeowner wants to minimise the cost of a retrofit but also minimise the carbon emissions of the property. These objectives are in opposition, as the homeowner could minimise cost by not installing a retrofit at all but in doing so they would not reduce emissions either. Similarly, minimising total emissions is likely to be one of the costliest interventions. It quickly becomes clear that multi-objective optimisation problems do not usually have a single solution, but rather, a set of trade-off solutions from which the decision-maker must make a selection based on their preferences between alternate objectives.

$$
\begin{aligned}
&\min\{f_1(x,y), f_2(x,y)\} \\
&\text{s.t.} \qquad\qquad\qquad g(x,y) \geq \epsilon \\
&\qquad\qquad\qquad\qquad\quad x \in \mathbb{R} \\
&\qquad\qquad\qquad\qquad\quad y \in \mathbb{R}
\end{aligned}
\tag{2.4}
$$

Equation 2.4: Generic multi-objective optimisation problem with two objectives, two decision variables and a single constraint.

Methods attempting to resolve multi-objective problems fall into two broad categories, a priori and a posteriori methods. A priori methods require knowledge of the decision maker's preferences at the time of optimisation. This allows the selection of a solution that best fits the preference provided. A priori methods often involve scalarising: using the preference information to form a single-objective function which can be optimised using traditional methods. Scalarising can be done through utility functions that map objective values directly to a preference measure, or through a simpler linear combination of the objective values (with or without weighting).

In contrast, a posteriori methods are those in which the preference of the decision maker is unknown at the time of optimisation. A posteriori refers to the ability to apply preference data after the optimisation is complete, allowing the stakeholder to make a decision from a range of alternate optimal solutions. As such a posteriori methods do not produce a single solution but rather a range of alternative solutions for stakeholders to choose from and so, in a sense, these methods do not solve optimisation problems but rather condense the decision space into more refined solutions. The standard metric for

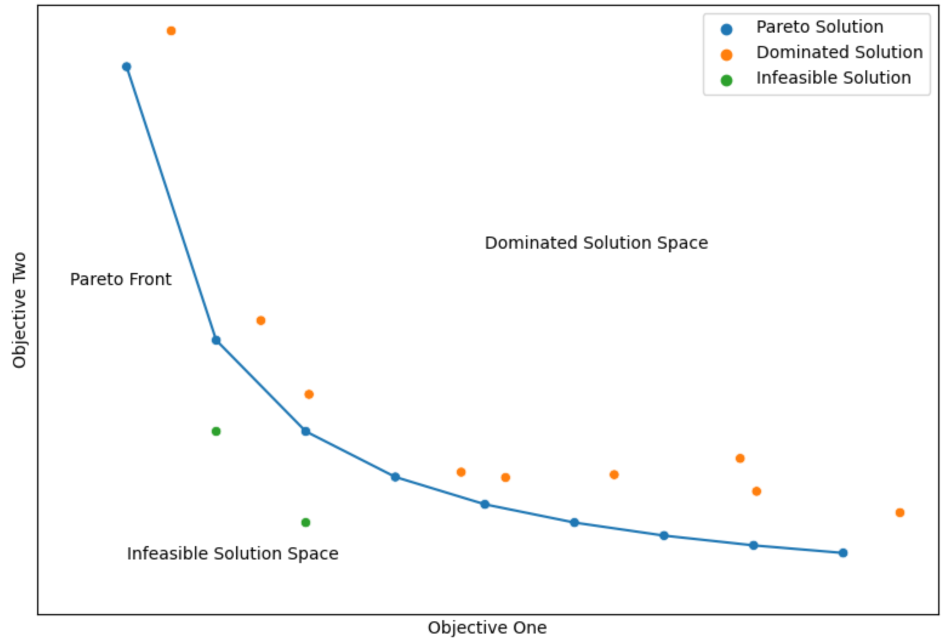

Figure 2.3: Pareto front shown for a generic optimisation with two minimising objectives.

which decisions are outputted from a posteriori methods is based on solution domination and Pareto efficiency.

Solution X is said to *dominate* solution Y if all objective values obtained from X are preferable to all objective values obtained from Y. A solution is Pareto optimal (or Pareto efficient) if no other solution dominates it, meaning that for all other solutions there exists at least one objective value in which the Pareto solution performs better. These solution types are shown graphically in Figure 2.3, which shows a generic optimisation with two minimising objectives. Any solutions found below the Pareto Front are infeasible, as feasible candidate solutions below a candidate front will themselves be Pareto efficient and thus form part of the front.

The reason that Pareto efficiency is vital to a posteriori methods is that, if a presented solution is Pareto optimal, then there exists a possible set of preferences that make that solution preferred. In contrast, if a solution is not Pareto optimal, then there exists at least one solution that would be preferable to a stakeholder regardless of their preferences. By attempting to seek Pareto optimal solutions, a posteriori methods attempt to list all of the solutions that may be preferred by a stakeholder while omitting all which would never be preferred. A set of Pareto efficient points is referred to as a Pareto Front, since, when objective values are plotted visually, Pareto solutions form a front between the dominated solution space and impossible-to-reach objective combinations.

There exist a range of methods for obtaining the Pareto Fronts for a priori methods. They can be broken down into two main categories. The first category of methods repeatedly scalarises the problem to generate an optimal solution for different possible preference sets, combining the found solutions into a front. These methods systematically divide the preference space and perform local Pareto optimisations at different points. Popular scalarisation methods include constraint-based methods, which systematically apply constrained optimisations to obtain constrained optimisation points, then recombine them (filtering out any dominated solutions) to form a front [35, 36]. The second category of optimisation methods attempts to generate and maintain an entire front of potential solutions at once, such as genetic algorithms. Algorithms that adapt GAs, such as the popular NSGA-II, attempt to maintain and improve on a whole front of solutions by ensuring non-dominated solutions are retained and guided towards optima, while dominated solutions are given lower priority [37]. This exploits the same optimisation advantages as a traditional GA and can generally obtain a front faster than repeat scalarization methods [38]. The downside of these solution-pool approaches is that, compared with repeat scalarization methods, they provide fewer guarantees of front diversity (as this is stochastic and controlled by the dynamic GA process).

2.3.3 Relationship between Machine Learning and Optimisation

Given that both machine learning methods and optimisation are prominent tools used in this thesis, it is worth considering how they relate to each other conceptually. Generally speaking, the training process of a supervised machine learning method is an attempt to optimise the parameters of a given model in order to minimise the loss. In this sense, the training process is an optimisation procedure where the decision variables are the model parameters and the loss is the optimal value. Training attempts to fit the parameters to best predict the desired outcome of the model. This does not, however, mean that the trained model is performing optimisation when predicting for unseen data (although it may be). Instead, predicting the target attribute, given the input attributes provided, which may differ greatly from the optimal prediction of a given stakeholder. This, therefore, depends both on what the model is trained to do (the loss function) and the objectives of the model user. A regression model trained to predict the heating demand for a building is optimally trained to make that prediction, but the stakeholder using the model is likely to desire a reduced energy expenditure. As such, the prediction is not an optimisation.

2.3.4 Building Energy Simulation

A more methodology-oriented discussion of building energy simulation is provided in Section 4.6.1, which focuses on the necessary practical components of the energy simulation software used in this thesis. Here we will discuss the fundamental concepts at play when simulating a building to ascertain the heating energy demand.

The primary energy use factors in buildings are heating, cooling, ventilation, lighting and plugs, and process loads [39]. It should be noted that not all of these concepts will be considered within the scope of this thesis, with factors like cooling omitted in all but discussion due to the current physical and socio-economic climate of the chosen case study. The primary focus of this research is on heating demand - the energy required to meet the heating schedule of a given building state. The primary building energy simulation tool used in this work is EnergyPlus. This software package has been funded by the U.S Department of Energy and managed by the U.S National Renewable Energy Laboratory [40]. It is one of the most feature-rich packages for simulation of individual buildings [39]. EnergyPlus is generally used for single-building energy simulations. Other tools, such as CitySim, are designed to simulate a larger area of buildings at once [41]. This is useful for instances where building interactions are significant, such as the shading of high profile buildings affecting lower profile buildings around it. Given the low-rise nature of the data set selected for this work, and the extra computational complexity of this scene-based simulation, the atomic simulation approach of EnergyPlus was used.

EnergyPlus simulations are defined in terms of an Input Data File (IDF), as well as an input data dictionary, and a weather file for the location of the building. The IDF contains the necessary components that define a building. The core components include thermal zones, areas for which the simulation monitors and calculates the thermal properties of surfaces that have thermal properties such as thermal resistance, reflectance, and capacitance. Surfaces are typically walls, doors and roofs, with windows making up a special case of glazed surfaces. In order to perform simulations, IDFs are constructed and altered based on building properties, a process we explain in more detail in Section 4.6.1.

The occupancy of a building, meaning the humans that exist inside the physical environment, can influence the energy demand in several ways. The most abstract method of capturing occupancy behaviour is through heating setpoints and scheduling. This captures the interactions between occupants and the control systems within the building in a predetermined manner. There are ways to capture more sophisticated occupancy behaviour using co-simulation. In EnergyPlus, this can be done using the functional mockup unit, which provides a runtime interface for external applications to observe and alter the built environment state, in line with simulation rules defined by the modeller [42]. This can be used for more dynamic behavioural elements, such as dynamic changes to a heating setpoint in response to thermal comfort, complex window behaviours of more precise control of electrical appliances. The approach outlined above, of setpoint and schedule based occupancy modelling, is used in this research due to the amount of data required to perform co-simulation accurately, as well as the performance constraints it imposes since it captures general heating control behaviour.

2.4 Decision Modelling and Social Simulation

2.4.1 Economic Models of Behaviour

The traditional microeconomic approach to modelling decisions begins with rational choice theory: modelling humans as rational agents who attempt to maximise their own self-interested goals through their action selection. These goals can be abstracted to the notion of utility, which is an attempt to capture the preferences of the rational economic agent numerically. The concept of utility is somewhat nebulous, encompassing a general sense of positive well-being from consumption that the agent prefers (increases utility) or a sense of disutility from consumption or actions the agent does not prefer (decreases utility). While utility cannot be directly measured, actions taken by individuals reveal preferences in given scenarios. These revealed preferences allow researchers to approximate the utility of given actions when compared with others. Analysis can also be performed by hypothesising utility functions based on real-world behaviour, survey data, or theoretical models, which can then be used to make behavioural predictions.

The notion that a rational agent will always act to maximise utility faces several theoretical hurdles. One issue with utility maximisation as a decision model is the notion of uncertainty. Consider a choice such as a selection between buying a lottery ticket or placing the equivalent value into a savings account; the most rational of those decisions is not immediately clear. If the lottery ticket was a winning ticket, then it would be rational to purchase it; whereas, a losing ticket would be an irrational purchase. This is overcome by the introduction of expected utility theory [43]. This expands on the standard rational-choice model by allowing for probabilistic outcomes, with utilities from given scenarios being weighted by the probability of that scenario occurring.

Under the theory of expected utility, it is possible to measure the degree of risk aversion by observing the choices made when decision makers are offered differing possibilities. Risk-averse individuals will select choices with higher probabilities of positive outcomes, while those with risk-seeking preferences may avoid scenarios with expected high positive outcomes due to their aversion to potential loss. The advantage of using risk aversion metrics is that it can be measured experimentally in laboratory settings, a topic which is well discussed by Harrison and Elisabet (2008) [44]. However, there is significant difficulty in determining risk attitudes accurately, as simple framing effects (such as the difference between a gain or a loss of the same amount) influence risk attitudes measurably [45]. A meta-analysis investigating the framing effects of 136 risk studies was performed by Kühberger (1998) [46]. While they found that framing effects were measurable, they were only small-to-moderate which suggests some transitivity of measured risk attitudes between frames. Individual models for predicting risk attitude have been used when modelling synthetic decision makers in agent-based models, despite the fact that getting accurate coefficients of risk aversion can be difficult even using participants in a lab [47, 48, 49]. However, measures of risk aversion purely derived from expected utility theory have been criticised due to the inconsistencies that appear from experimental data, with some suggesting alternative models that are less rigidly tied to utility [50].

2.4.2 Public Goods Games

Public goods games are highly relevant to climate modelling. This field of experimental and behavioural economics is concerned with what economists consider public goods, which have two key characteristics:

1. Non-rivalrous. Once produced, the consumption of the good does not reduce the availability of the good to others.
2. Non-excludable. Once produced, it is not possible or feasible to selectively deny access to the good.

Examples of public goods include national defense and law enforcement, as these systems being in place benefit the entire population. An unpolluted environment is also a public good, as it benefits everyone without detriment to others' ability to do so (non-rivalrous) and which cannot be selectively denied to individuals (non-excludable). In traditional economic theory, public goods result in market failure through the free rider problem. This occurs because an individual reaps the benefit of the public good regardless of their contribution towards it, which is unpreventable due to its non-excludability. Considering an environmental issue like reducing emissions, a household benefits from the effort that all other households put into reducing their emissions, regardless of their own contribution to the initiative. This fact, coupled with the reality that the marginal benefit that a single household can contribute to an emissions reduction initiative is small relative to total emissions, can result in a free riding problem where households do not engage with the issue of emission reduction as they believe their contribution makes no difference.

The results of the public goods games carried out by economists appear to confirm this market failure. In the most basic version of the game, each participant (i) is endowed with an initial sum of money ($epsilon_i$). They are able to contribute a proportion (c_i) of this money towards a public good P using their endowment, which 4 other players may also contribute to. This results on a public good spending of $P = \sum_{i=1}^{5} c_i$. The sum received by participants at the end of the experiments is $(e_i - c_i) + MP$. For all $0.2 < M < 1$ the socially optimal solution is for full contribution to the public good, while the profit maximising Nash equilibrium is no contribution to the public good. A typical game to model a phenomenon such as environmental emissions would be a repeated public goods game, whereby participants can observe the contributions made by others between rounds before determining their future contributions. Experimental results show that in single shot games there is a positive contribution to the public good, but that in repeat games participants converge towards free riding [51].

The modelling of altruistic decisions, such as contribution to public goods, is something that appears at odds with the traditional notion of a self-interested rational agent considered in classic economics. As such, altruism has been studied by both experimental and behavioural economics in an attempt to model its behaviour pattern within a rational framework. To comply with general utility theory, an agent who is witnessed to forego personal consumption to perform an act of altruism must achieve greater utility from the act of altruism than they do from the foregone consumption. The 'warm glow' theory, for example, suggests that positive utility is gained when doing something perceived as virtuous and this has been linked to public goods contributions [52]. There is also evidence of public good contribution on the grounds of conditional cooperation, in which participants are likely to contribute to a public good if they observe or perceive others as doing so [53].

2.4.3 Investment Decision Modelling

When considering a retrofit as a form of investment, there is a wide range of tools that can be used to evaluate them. The investment decisions in the scope of this research, namely retrofits installed on buildings, are generally categorised by a large upfront cost which is offset by future energy bill savings for the lifecycle of the retrofit. One simple but highly effective investment decision metric to model this formulation is Net Present Value (NPV). NPV, as shown in Equation 2.5, is the current value of a flow of future returns (R_t at time t), from initial investment C_0, and discounted in a compounding manner by a rate i. As such, near payments are valued more highly than future payments. The discount rate applied to calculate NPV can be conceptualised in several different ways. In financial calculations, it usually represents the value of interest that near payments are able to earn compared with temporally distant ones, capturing some of the opportunity cost of any value tied up in future payments. It can also be considered as an offset for expected future inflation, as the spending power of payments will be reduced when inflation occurs.

NPV is established as a viable decision tool in the investment and adoption literature [54, 55, 56]. One advantage of NPV is that its unit is present-value currency, making it easy for decision makers to understand. This is a valuable property when constructing preference calculations as it allows comparison to other preferences if they can be expressed financially, as we do in Chapter 6 with consumer's willingness to pay. In general, a rational agent wishing to maximise financial gains will make any decision with a positive NPV, provided the discounting rate is accurately calculated and the decision does not impose opportunity costs. In reality, opportunity costs are often present, with some financially beneficial decisions precluding other, more beneficial ones.

$$NPV = C_0 - \sum_{t=0}^{n} \frac{R_t}{(1+i)^t} \tag{2.5}$$

Equation 2.5: Net Present Value for an investment of C_0, returns of R at time t and a discount rate of i.

One valid criticism of NPV as an investment decision metric is the inability to accurately capture opportunity costs. NPV calculations do not account well for factors like future changes in technology, thus failing to account for some possible opportunity costs that can be captured by inaction [57]. Consider a simple scenario in which a positive NPV retrofit solution is installed, only for a new technology to be released during the lifespan of the installation. Even though the initial installation was of positive NPV, a rational agent without perfect foresight would perform worse than an irrational one who only acted after the technology came to market.

While NPV is a good decision metric, there are some other good metrics to consider when evaluating an investment decision. Return On Investment (ROI), shown in Equation

2.6, represents the return of a given investment as a proportion of the costs incurred. The return, R_l should be expressed in present value, with appropriate discounting, and accounting for ongoing costs. This form of representation may be preferable when initial funds are constrained, as a decision maker may opt for a lower-cost investment with a higher rate of return, even if the NPV is lower. The conceptual importance of this metric depends on how well opportunity costs are captured in the NPV, as a rational agent should still prioritise the highest NPV alternative, even with lower returns, provided the NPV captures the liquidity costs and discounting of the investment appropriately. Consider, for example, a small investment with a lifecycle cost of £1000 and lifecycle returns of £1100, compared with a larger and mutually exclusive investment of £10,000 returning £10,500. The ROI of the small investment obtains an ROI of 10% and an NPV of £100, compared with 5% and £500 of the larger investment. Provided the NPV calculation includes the cost of borrowing (discounted correctly) and the larger investment bears no additional risk, then the larger investment is preferable despite its lower ROI.

$$ROI = \frac{R_l - C_l}{C_l} \tag{2.6}$$

Equation 2.6: Return on investment with a lifecycle cost of C_l and lifecycle return of R_l.

Chapter 3

Related Works

This chapter will discuss the literature most closely relating to this research, identifying and critiquing the existing tools for analysis of energy retrofits in urban building stocks. Section 3.1 discusses the methods for evaluating the performance and optimising of domestic dwellings in isolation. This is then extended to the methods for evaluating and optimising the retrofits of entire building stocks in Section 3.2. Section 3.2.2 includes some discussion of the decision models required to predict the actions of those making retrofit choices. This is explored in more detail in Section 3.3 which examines the use of Willingness to Pay (WTP) for contributions in environmental decision models. Finally, a discussion of the research gap identified across the literature can be found in Section 3.4.

3.1 Energy Performance Modelling and Optimisation

There exist a variety of tools used by researchers and practitioners in evaluating the energy performance of domestic dwellings. Some of these tools have been integrated into optimisation procedures which allow for the evaluation and efficient improvement of performance. This section will provide an overview of these methods. The scope of the methods in this section will generally be limited to individual dwellings, with a discussion of methods for evaluating and optimising entire building stocks performed in Section 3.2.

3.1.1 Building Performance Evaluation

Building Performance Evaluation and Certification (BPEC) tools are a category of procedures for determining and classifying the performance of dwellings' energy demand and carbon emissions. The most common BPEC tool used and required by government departments for energy and environmental assessment of dwellings in the UK is the Standard Energy Procedure (SAP) [58]. This methodology is used in official procedures such as the generation of energy performance certificates, calculating emissions for stamp duty exemptions, and to ensure building regulation compliance. While we will focus mostly on the UK's SAP methodology, as it relates most closely to the building stocks considered, different countries operate using different BPEC procedures. The United States uses a home energy score system, which integrates more simulation tools (such as EnergyPlus) rather than relying on analytical modelling techniques [59].

The SAP model is analytical in nature and was designed by BREgroup, a certification and standards body [60]. The documentation of the most recent model, updated in 2012 can be found on their website [61]. The methodology is designed as a standard for easy comparison between heterogeneous dwellings and, as such, standardised assumptions of occupancy are made. This is desirable for processes such as building regulations, which may be required before any occupancy information is available.

One flaw identified in the traditional SAP methodology was the data requirements. While uses, such as regulatory compliance at the planning stage, are likely to have access to complete building data, stock assessments may be working with significantly more limited inputs. In response to this problem, the Reduced Standard Energy Procedure (RdSAP) was developed, reducing the data requirements and thus the level of detail of the assessments, making this methodology more suitable for modelling existing buildings.

In many ways, analytical BPEC tools such as SAP are well suited to building stock modelling of the sort performed in this research. These mathematical models are generally computationally cheap to perform, compared with full simulation procedures. The models, particularly those such as RdSAP which are designed for existing buildings, require relatively little data so as to allow assessment of existing buildings without access to architectural designs or details that are unavailable to a casual assessor.

However, while the SAP and other standardised analytical models are generally computationally cheap compared with full simulation models, their fixed form comes at the cost of reduced flexibility. The types of model inputs for the analytical models are fixed during construction with assumptions, such as occupancy details, set as defaults. This is a desirable property for standardised models, as it minimises the information required about these inputs and allows for a fair comparison between the intrinsic properties of dwellings. Still, this makes the models inflexible and introduces known error from under-specification. It can also mean some important components are missing from the analysis. This is addressed by Kelly et al. (2012) [62], who critiqued the prior version of the SAP in their recommendations for the expansion of the previous SAP methodology. This was due to missing components such as fuel sources with varying levels of decarbonisation. It could be difficult, for example, to model a decarbonising grid scenario using the SAP methodology. In contrast, simulation-based assessment relies on fewer assumptions, allowing the modeller to specify input variables of interest while still permitting the use of defaults.

The relative inflexibility of constructed analytical BPEC procedures can be detrimental to some applications, such as building stock modelling, when the initial model was designed for a different purpose. While models like the RdSAP were specifically designed to require less data than the SAP in order to model existing stock, the model design is based on the assumption that the assessor would be able to access and assess the property from the inside, measuring the areas of various architectural properties on site. This assumed data collection method is suitable for several intended purposes, such as evaluating the performance of a rental property but may not be well suited to simulation and optimisation of an entire building stock in which researchers are unable to enter individual dwellings for data collection. In this instance, simulation-driven surrogate models give the researcher more flexibility, as unavailable input variables can be avoided during model construction. This may result in an under-specified model, compared with the constructed analytical alternative, but allows simulation in lieu of certain inputs.

While the analytical BPEC methodologies such as the SAP and the RdSAP may lack the flexibility desired for building stock analysis and optimisation, the models themselves may be useful in this research. The SAP can be considered a grey-box modelling method, as they are less transparent than full energy simulation techniques but the effect of input changes is clear from the mathematical form of the model. As such, the relative importance of given inputs can be compared with data-driven Surrogate Energy Performance Models (SEPMs) as a form of model verification. If, for example, the most significant input variables from the SAP methodology are shown as insignificant in a trained SEPM, it may indicate issues with the data-driven model. These models may also be useful for guiding and justifying design decisions when creating data-driven SEPMs. The SAP methodology, for example, uses separate thermal zones for living and non-living areas, which is determined to be a good representative approximation for domestic dwellings [60]. It may therefore be desirable to replicate this property as a standard assumption when designing a simulation framework for a SEPM.

It is worth considering, conceptually, the difference between constructed analytical models such as the SAP and the data-driven SEPMs of the type discussed in Section 3.1.2. While the SAP is designed for a particular purpose by experts, the model design is still driven and calibrated by using both real-world and simulated data. To use the modelling terminology established in Section 2.1: the SAP is a general model that, when filled with building-specific input data, can be used to simulate the target system (the individual dwelling in question). In this sense, it fulfills the same purpose of capturing the energy demand of a given domestic dwelling, as the SEPMs discussed in Section 3.1.2 and used throughout this research.

3.1.2 Surrogate Energy Performance Models

Surrogate modelling, also known as meta-modelling or response surface modelling, is the use of a faster but less accurate model to replace a slow process [63, 64]. This process can be understood as a model for which the target system is another model, one which is infeasible to construct or run. Surrogate modelling is generally done for the speed and resource benefits of the surrogate model over the original simulation model but there may also be advantages in the reduction of data requirements compared with a more rigid target model. While the target system of the surrogate model will be another model, the ultimate target system is therefore the same as the original model. A surrogate model targeting a building energy simulation model is, therefore, ultimately targeting the real-world building that the researcher desires to simulate. This method does introduce an extra source of error between the surrogate model and the ultimate target system. However, one advantage of surrogate modelling is the ability to run the target model to understand the exact magnitude and distribution of the introduced error, a process of validation that may not be possible against the ultimate target system.

Building energy simulation, the attempt to capture the underlying energy processes affecting a building in a high level of detail, as described in Section 2.3.4, can be computationally infeasible in scenarios such as optimisation loops where many repetitions are required. This is exacerbated in bottom-up building stock analysis, as these optimisations are required for a large number of buildings in a stock. The use of Surrogate Energy Performance Models (SEPMs) to reduce the computational cost of simulation stages is a common method in sustainable building design [65]. SEPMs are surrogate models which attempt to capture the behaviour of models from which they are trained. Due to the availability of the original model, and therefore the ability to generate the required training data, supervised learning techniques are well suited to these tasks. A background discussion of this family of techniques can be found in Section 2.2.

The importance of using a surrogate model for this process lies in speed. The optimisation process is slow, for example, Aijazi (2017) [66] notes that the 7000 iterations required to optimise each building in their model would have required approximately five days of computing time (with approximately 1 minute per simulation). Using a trained surrogate model framework the iterations were orders of magnitude quicker (just 0.0006 seconds per cycle), allowing for an optimisation time of just 4 minutes. In later work, the author's framework was expanded, allowing for calibration of urban building energy models 500 times faster than using non-surrogate methods with more reliable degrees of accuracy [67].

Surrogate modelling can be performed with a variety of different ML techniques. Tseranidis, Brown, and Mueller (2018) [68] compared six surrogate modelling techniques. The artificial neural network based surrogate model was found to minimise error across both case studies and performance metrics. They also found Kriging regression performed well, with random forest and radial basis function models under-performing by the error metrics used. The largest scale SEPM we have seen to date was presented by Edwards et al. (2017) [69] who trained a Deep Neural Network (DNN) surrogate using a big data approach, achieving high levels of accuracy at hourly precision with errors of less than 5%. This level of precision is often unnecessary, with total energy demand calculations being sufficient to analyse the impact of most retrofit installations, which they were able to calculate with errors of just 0.07%, at a greatly reduced computational time compared with traditional energy simulation.

One of the key details in SEMP design is the selection and relative weighting of parameters and model construction. Tian et al. (2015) [70] highlight this when they performed an exploratory study of university campus buildings to ascertain which key factors are correlated to energy usage. They implemented both linear and non-parametric regression models, finding that in this case the linear model had the best predictive accuracy. The variables they investigated included attributes such as the area and height of the building, the amount of glazing on each side of the building, and the duration and density of building occupation. They found the most significant predictive factors in their analysis were wall insulation, occupant density, and roof insulation.

When trying to improve the speed-to-accuracy trade-off inherent in SEPMs, Westermann and Evins (2021) [71] used a hybrid approach by training a SEPM with variance inference. As predictions come paired with a predicted uncertainty of the prediction itself, they were able to perform slower, high-fidelity simulations only in instances where the model predictions were expected to render the most error. In doing so, they reduced the error by up to 30% on the highest error samples. This is an example of integration between statistical and engineering models for energy performance evaluation.

There is relatively little literature attempting to integrate intelligent human behaviour into SEPMs, although there have been some attempts that focus on setpoint based evaluation. Wate et al. (2020) [72] demonstrate a novel surrogate model framework that emulates a stochastic building performance simulator, trained using EnergyPlus and co-simulated with the Multi-Agent Stochastic Simulation platform NoMASS, to simulate occupants' behaviours. The authors build a pair of surrogate models for the mean and variance of annual energy use and utilise these to decompose the impact of different sources of uncertainty. They found that in a simple mono-zone office building, the uncertainty in insulation thickness on heating demand dominated stochastic elements of human behaviour, but that the opposite was true in predicting cooling demand, where occupants' stochasticity dominated.

3.1.3 Optimisation using Surrogate Models

A common use of SEPMs is integration with an optimisation method. Due to the potentially large number of function calls required for optimisation, the use of a surrogate greatly reduces the computational cost. This is the most common method of optimisation using SEPMs, as well as the most common use of SEPMs found in the literature.

This method of surrogate modelling based optimisation is agnostic to the specific surrogate modelling technique, optimisation algorithm, and objective values. For example, Magnier and Haghighat (2010) [73] used an ANN-based surrogate model paired with a multi-objective GA for optimisation. The flexibility of DNNs combined has made them a popular surrogate modelling and optimisation technique for objective value estimation, with another good example by Ascione et al. (2017) [74] who evaluate cost, energy savings and thermal discomfort of solutions using trained DNNs. They designed the system to be generic and applicable to any building type, although it was not targeted towards whole stock modelling.

Prada et al. (2018) [75] provide an analysis of the performance of different surrogate modelling techniques in the context of optimisation. They used a GA with alternate surrogate models, comparing results to a brute force optimisation to determine the efficacy of alternative methods. They confirmed the practice as an acceptable way of optimising, with the majority of Pareto solutions identified by simulating only 3%-8%

of the solution space for training.

Waibel et al. (2019) [76] performed an extensive investigation of the optimisation techniques applied to building energy optimisation. They did not find any algorithms that outperformed the others on all fronts, with standard meta-heuristics such as genetic algorithms performing well. They also investigated some model-based methods of the type discussed in Section 3.1.5. While there are minor performance differences for the different optimisation methods, varying depending on the computational budget and specific problem specification, all of the black-box optimisation methods tested performed competitively, although parameter tuning was found to be of significant importance. Background information surrounding some of the optimisation methods considered can be found in Section 2.3.

Sharif and Hammad (2019) [77] use an adapted version of the optimisation with surrogate modelling method by generating a sample of Pareto fronts using traditional simulation-based multi-objective optimisation techniques. They use a three objective function method by training two ANNs: one which pairs the energy consumption against lifecycle cost, and another against lifecycle assessment. This demonstrates a capacity for extending the technique into further dimensions, although the scalability of this method as the number of objectives is increased is not clear. Their approach is novel, using only non-dominated solutions to train the data set, reducing the generality of the model in order to increase efficiency. This trade-off is also true in feature selection, as the model is designed to be used on a single building per model rather than a stock-based approach. While more expensive computationally than training an energy-focused SEPM, this method allows for more nuanced objective values such as building specific lifecycle analysis which may not be as simple to model as energy performance.

The popularity of ML and surrogate models in energy performance modelling has encouraged recent attempts to provide open source tooling. Westermann et al. (2021) [78] provide and document a toolset developed in Python that contains a pipeline for the development of SEPMs. The platform BESOS (Building and Energy Simulation, Optimization and Surrogate Modelling) utilises EnergyPlus for building simulations, ML libraries for surrogate model formulation, and meta-heuristic approaches (such as genetic algorithms) for optimisation. While these elements have been used in other work, the combination of these tools into an open source platform, as well as the data structures and design principles implemented for users, mitigates some of the domain knowledge required to perform this research.

A general review of the use of surrogate modelling in sustainable building design can be found in Westermann and Evins (2021) [79]. They assessed 57 surrogate modelling papers, finding that the most popular simulation tool was EnergyPlus, which is used to perform the initial energy simulations for SEPM training. While they identified eighteen uses of surrogate models for optimisation, they found most focused on optimisation of single buildings. This often involved training surrogate models with local scope; they cannot be generalised to other buildings or building types. While linear regression models were most popular across all the studies reviewed, ANNs were used more often for the purpose of optimisation. They only highlight one study which focuses on a generalised built-building surrogate model which could be suitable for optimisation of multiple buildings [80].

The most generalised and relevant use of surrogate models for retrofit analysis is presented by Ascione et al. (2017) [80]. The model was generalised using geometric summary features, such as the floor area, number of floors, and glazing ratios. An ANN-based model was chosen due to the increased generalisability granted compared with, for example, linear regression models. Nonetheless, the researchers found that the highest level of accuracy was kept by training separate models for the existing stock, compared with the retrofit stock. Retrofit analysis options used in the study were kept relatively simple, with specific retrofit decision rules applied to the entire stock, compared with an optimisation procedure. This paper uses a surrogate model for bottom-up Housing Stock Energy Model (HSEM) of the type discussed in Section 3.2.

3.1.4 Surrogate Optimisation

There have been some recent attempts at developing pre-trained predictors of near-optimal retrofit solutions, referred to in this thesis as surrogate optimisation. The target system for these models is not just the energy performance of the building, as is the case in SEPM, but also the optimisation and decision process that is used in the selection of the final energy retrofit. This is done to further reduce the computational cost of obtaining a retrofit solution for a given stakeholder's preferences. While standard optimisation procedures used in conjunction with SEPMs are suitable for the evaluation of individual buildings, optimising an entire stock of buildings controlled by heterogeneous stakeholders, as discussed in Section 3.2, may be too computationally expensive. Surrogate Optimisation could be a suitable technique for this task, although very little literature could be found with this application.

A single objective whole-city surrogate optimisation was attempted by Hey et al. (2020) [81][1], who trained a SO on a sample of buildings in Nottingham city in order to predict

[1] This work was published as part of the research presented in this thesis.

Net Present Value (NPV)-optimised solutions for the remaining stock. A financially viable solution was identified for 16.7% of the whole building stock compared with 19.2% of the training sample, suggesting 87% of buildings with a viable solution were identified. Predicted solutions performed 11% worse than the sample solutions, making the significant computational savings of the approach a trade-off with solution quality.

In 2021, Thrampoulidis et al. (2021) [82] presented a hybrid classification and regression model to predict near-optimal retrofits using cost and emissions objectives. Their classification model used preset wall, window, and roof insulation settings to perform binary classification, reducing the complexity of the problem compared with a regression or multiple-classification technique. However, they focused on a wider range of energy supply methods, including renewable sources and energy storage. Their method generates a fixed-length Pareto front by first generating a solution that maximises a single objective value, then making sequential calls to a model in which the output of the previous stage is used as input for the next. This method is efficient for generating a front of known size but obfuscates the underlying trade-off required to move from one point on the front to the next. Their work found surrogate optimisation to be a convenient balance between computational cost and accuracy, as well as potentially being more accessible for non-experts to use due to the diminished data requirements of the trained model compared with traditional retrofit optimisation approaches.

3.1.5 Model-Based Methods for Optimisation

There exist model-based optimisation methods which attempt to construct a surrogate model of the fitness landscape during the optimisation process, in order to converge more efficiently to the optimum solution [83]. There is active research in the area of combinatorial optimisation, which attempts to apply deep learning to capture patterns in a given solution space to perform optimisations [84]. These methods are quite promising and have been proposed as alternatives to more common metaheuristic approaches such as GA optimisation [85]. There is some concern however that while they converge on solutions quickly, they are more susceptible to getting stuck in local minima, suggesting insufficient diversification [76]. While deploying a similar concept to surrogate optimisation discussed in Section 3.1.4, these algorithms are trained in real-time to optimise a specific objective function, in contrast to a pre-trained model, to achieve fast optimisation predictions on genetic buildings for heterogeneous stakeholders from a data set by training using a metaheuristic approach.

3.2 Housing Stock Energy Models

Housing Stock Energy Models (HSEMs) represent an attempt to model the energy behaviour of an entire stock of buildings, which may range from an entire national stock to smaller regions, cities or districts [86, 87, 88]. Building stock analysis can be approached from either a top-down, bottom-up or a hybrid perspective [86, 89, 87, 90]. These general modelling paradigms are discussed in Section 2.1.3.

Swan and Ugursal (2009) [86] review and classify a range of national and regional HSEM techniques. They identify the advantages and drawbacks of top-down and bottom-up modelling methodologies. They also subdivide bottom-up techniques into statistical and engineering focused methodologies. The statistical approaches involve modelling individual households/archetypes using statistical modelling techniques derived from historical data [91]. While these techniques give more granularity than top-down approaches, they still suffer from a reliance on existing historical energy consumption data and may limit occupancy behaviour modelling to that which has been recorded and observed. In contrast, bottom-up engineering approaches allow simulation of new technologies and different models of occupancy behaviour, yet they come with their own challenges, in the form of building specific data requirements and high computational cost [86, p. 1833].

Lee and Yao (2013) [92] discuss top-down and bottom-up approaches in the context of new energy technology adoption. They advocate a bottom-up, particularly agent-based model due to the higher resolution and level of detail that can be included in these models. They constructed a demonstrative ABM of retrofits, although retrofit options were limited to boolean representations: each measure was either installed or not installed. This led to only 32 different combinations of retrofit installation, making optimisation of each dwelling trivial. As such, when they highlight the drawbacks of the agent-based approach, they focused mostly on relating it to the availability of data rather than the computational cost. However, this boolean representation of installation measures resulted in limited details about the dwellings at hand, as the nature of the retrofits to be installed were so simplified.

Hall and Buckley (2016) [87] reviewed and classify general energy system models found in the literature between 2008 and 2016. Their categorisation procedure indicated that while bottom-up models were common for energy supply modelling, the energy demand from buildings' models they identified generally relied on top-down techniques. Although their review was generally focused more on broader energy system models, making a bottom-up analysis less feasible (as the built environment represents only a small component of the system studied).

One popular method used to create HSEMs is to pair analytical physical models with statistically assigned housing stock data, scaling up results to represent the whole stock. An example of such a model is the Cambridge Housing Model, which combines the SAP methodology discussed in Section 3.1.1 with building data from the English housing survey [93]. A sample of dwellings is generated using the survey data, cleaned and evaluated using the SAP methodology and then extrapolated to the remaining population of dwellings. This method benefits from the speed advantages of analytical models but suffers from the combined inflexibility of the expert-derived physical models as well as the drawbacks of sampling and extrapolation [88].

Sousa et al. (2017) [88] reviewed 29 national scales and found that the tools being utilised by researchers are "limited in their scope and employ simplistic models that limit their utility". To combat this, they developed EnHub [94], a transparent and modular platform for the analysis of national housing stocks. This, initially UK-focused solution, uses a hybrid of top-down and bottom-up methodologies by both clustering and sampling the building stock to form archetypes, before extrapolating up to the entire stock. They first identified a set of archetypes by data mining and reducing samples from the English Housing Survey (EHS) data set. A total of 64 geometric archetypes were identified, which were further supplemented with semantic attributes such as age and tenure to ensure the most statistically relevant properties extracted from the data were captured in the samples. Using a Latin Hyper-Cube sampling technique on these identified variable combinations, the data set was reduced from the original 14,951 to 1,016 total archetypes. Next, the researchers developed a modular simulation framework to dynamically generate EnergyPlus Input Data Files (IDFs). Their simulation framework followed the BREDEM structure used in the SAP methodology and discussed in Section 3.1.1. Analysis of the national building stock was extrapolated from the simulated performance of the archetypes. The researchers performed some basic retrofit scenarios by considering the uptake of a relatively limited set of retrofit options. They included a set of 8 retrofit policies and applied them in two behavioural scenarios: A perfect uptake scenario that captured the saving potential in the stock, and a conditional uptake scenario that attempted to emulate some investment decision components by controlling for the technical and economic feasibility of a given retrofit.

Another example of a bottom-up approach based on sampling and upscaling is presented by Amstalden et al. (2007) [95]. They constructed a bottom-up model of the Swiss residential building stock to investigate the retrofit potential within the stock. They used a discounted cash flow model to calculate the financial returns of retrofits and to act as a single-objective decision model when determining their application. They only included four pre-selected WHRS combinations for analysis rather than performing an optimisation procedure on individual dwellings. While this made their method computationally feasible, the choice of retrofit for each dwelling was very limited.

There have been relatively few uses of Surrogate Energy Performance Models (SEPMs), of the type discussed in Section 3.1.2, for bottom-up building stock modelling. Melo et al. (2014) [96] performed a feasibility study examining the use of an ANN for surrogate modelling of the urban scale building stock in Florianópolis, Brazil. Their SEPM was trained using EnergyPlus simulations. Their analysis considered the effect of optimising building-shell characteristics but analysed retrofits purely from an energy performance perspective and did not come across the scaling issues that occur when performing full optimisations on each residence. Mastucci et al. (2014) [97] use a linear regression model to evaluate the energy demand impact of four pre-selected retrofit bundles on the city of Rotterdam. This limited choice of WHRS combinations again demonstrates the computational complexity of performing retrofit optimisation at this scale, with the authors instead selecting a small number of fixed measures to be applied.

3.2.1 Urban Building Energy Modelling

Urban Building Energy Models (UBEMs) are HSEMs focused on the urban scale. Sola et al. (2018) [98] reviewed the UBEM tools available in 2018 with a focus on bottom-up modelling approaches. Simulation tools such as EnergyPlus are used for modular, single buildings or small district simulations in work such as Huber's (2011) [99]. Although this seems to have only been done when examining smaller regions or in hybrid methodologies which extrapolate from archetypes. This is generally modelled by extrapolating archetypes to an entire city. When considering the impact of retrofits, models rely on limited and pre-selected WHRS combinations rather than performing an optimisation procedure [98, 97, 95]. This is due to the computational cost of WHRS optimisation at this scale of housing stock.

There are some simulation modelling tools specifically designed to simulate building energy usage in urban environments. To represent the geometric data, City GML (Geography Markup Language) was proposed [100]. This allows for a common language in the development and sharing of city models [101]. City GML is also used as an input to urban scale energy simulation tools such as City Sim [41]. Work has been done to generate City GML from open public data sets to allow for dynamic urban stock modelling [102, 103].

Issermann et al. demonstrate an UBEM in which templates are used and manipulated to create EnergyPlus input files based on available building data [104]. This bottom-up modelling approach is supported by a functional mock-up interface to allow for urban scale simulation using EnergyPlus, which is designed for smaller scale simulation procedures. However, the energy model was limited to a residential block of 5,736 buildings, likely due to the computational constraints found with bottom-up simulation modelling, making this methodology alone unsuitable for optimisation of large scale stocks.

While the simulation of individual buildings can be feasible in UBEM design, the use of archetypes and extrapolation is also common. Pasichnyi et al. (2019) [105] discuss and present a method of archetype formation in the urban setting, using Stockholm as a case study. The use of archetypes for upscaling energy performance calculations allowed the assessment of seven retrofit packages on a stock of 5,532 dwellings. However, their energy assessment was performed on the basis of energy performance certificates rather than physical models, limiting the fidelity of outputs to those which are available in the data set. The evaluation also relied on the use of pre-selected retrofit packages rather than performing a WHRS procedure for the selected dwellings.

3.2.2 Retrofit Decision Models

Modelling the performance of a building with a given retrofit is only one aspect of retrofit modelling. It is necessary to select a retrofit solution from all the possible retrofits that could be performed. Much of the literature surrounding domestic retrofit optimisation is concerned with optimisation of the retrofit choice against easily quantifiable objectives such as Return On Investment (ROI), thermal comfort, or environmental impact. While it is possible to optimise for a single objective, it is more usual to perform a multi-objective optimisation of the type discussed in Section 2.3.2. Often single building retrofits are presented a posteriori as a Pareto front of viable solutions that can be offered to a decision maker to judge. However, the presentation of a range of equally viable solutions for each residence is not always desirable, as it does not allow researchers to make predictions of the actual behaviour of a building or building stock.

To solve the problem of retrofit selection, researchers can use a decision model. The target system that decision models attempt to emulate is that of the real world decision maker that would decide which, if any, retrofit should be installed on a given dwelling. It should be noted that not all research that does this will refer to a decision model in those terms but, for the purposes of evaluating the literature, any method which selects a single retrofit solution in the way that a real-world decision maker would, will be referred to as such.

While there are decision models in use by researchers, it should also be noted that another related set of methods are often referred to as decision models, although they are more often referred to as decision support frameworks, in the literature. These are prescriptive frameworks are designed as tools to aid real decision makers in selecting retrofit solutions. Examples of decision support frameworks can be found in Menassa and Baer (2014) [106], Hong et al. (2014) [107], Pazouki, Rezaie and Bozorgi-Amiri (2021) [108], and Duah and Syal (2016) [109]. These prescriptive methods attempt to influence or improve the decisions made rather than to emulate how a real decision maker may, in the absence of the framework, make a decision. It is possible for a prescriptive method of this type to be used descriptively, although in doing so the researcher makes the assumption that decision makers will be using the prescriptive model, or a close equivalent, in guiding their decision.

3.2.2.a Optimisation Based Decision Models

One method of retrofit decision modelling is to perform an optimisation. While a posteriori multi-objective optimisation methods are unsuitable, as they require a decision maker to select from the solution range. These methods have been used as part of decision support frameworks, as they provide stakeholders with an intelligently pruned selection of solutions along a broad range of retrofit attributes [108]. It would be possible to pair an a posteriori optimisation method with a second decision model, to first find a set of Pareto solutions and then select them, but no attempts to apply this method to building stock energy retrofits were found. While a posteriori multi-objective optimisation techniques alone do not solve the decision problem, it is possible to use single objective optimisation and a priori methods to obtain a solution.

As discussed in Section 2.3.2, multi-objective optimisation methods can be scalarised by providing preference information that can be used to balance alternative objectives into a single value, allowing for single-objective optimisation to be performed. When scalarising using an objective function, the function and its parameters can be conceptualised as the model of the decision maker, as it captures their preferences. These scalarised, objective-function based decision models create utility-based agents, discussed in Section 2.1.4. Asymmetric agents, with different utility function parameters, are suitable to represent the preferences of a wide range of decision makers, as they may be found in domestic building stock.

A similar method of scalarising the optimisation problem is to select a criterion that already compounds several key criteria, such as the lifecycle cost or NPV. Chiara and

Bragolusi (2018) [110] performed a systematic review of the valuation techniques used in retrofit literature. Common techniques included lifecycle cost, NPV, and lifecycle greenhouse gas emissions. The values selected demonstrate one of the large drawbacks of scalarisation techniques for decision modelling. Scalarisation can be done with minimal knowledge of stakeholder preferences when objectives are of the same type, such as financial values which can be converted into present-financial value using a discount rate. However, objectives of a different type, such as carbon emissions and financial return, require preference information for scalarisation with an objective function.

There has also been some criticism of using scalarisation techniques for retrofit decision making. This stems from the significant uncertainties present at the decision point of retrofit choice and the conceptual confusions caused by atomising a multifaceted problem into a single unit of value. Some advocate for less emphasis on numerical decision support tools and greater understanding of the cognitive decision making process when presenting retrofit options [111].

3.2.2.b Rule Based Decision Models

One solution, particularly popular in HSEMs, where the evaluation of individual dwellings is too computationally expensive, is to use rule based decision models. These can be considered analogous to reflex agents of the type considered in Section 2.1.4, acting in a predetermined way to specific environmental conditions. An example of rule based retrofit decision models, applied to housing stock, can be found in Sousa et al. (2018) [94]. Researchers implemented two rule based decision models; in the perfect uptake scenario, all dwellings were retrofitted into the most efficient form. In the conditional uptake scenario, retrofits rules were constrained by practical and financial limitations. The use of rule based analysis is common as a method of evaluating retrofit potential when computational constraints prevent optimisation, and is found in other work discussed in this chapter [112, 105].

When obtaining an Energy Performance Certificate (EPC) in the UK, a recommendation report is provided alongside the analysis of existing performance [113]. This report follows a simple rule based decision process wherein properties lacking a certain feature (double glazing, low energy lighting, etc.) are recommended that feature if the dwelling is suitable. In this sense, the recommendations could be considered more like minimum standards to which dwellings hope to conform to. Given that EPCs are produced using an RdSAP energy performance model, these recommendations also include an estimated payback period based on estimates of energy savings. In this sense, the inclusion of only cost effective measures extends this decision model to approximate a purely rational one, with only positive economic returns considered.

Khayatian et al. (2017) [114] used a classification model for label prediction of a ML model. However, training is based on observations from an energy certificate-based data set, resulting in a rule based decision model. In this sense, the ML model attempts to generalise the existing rule based decision model to uncertified dwellings. While this is useful in obtaining decision predictions, the quality of the predictions is capped by the underlying decision model used on the training set, which is generally low resolution and based on brief assessments.

3.2.2.c Analytical Decision Models

Analytical decision models are generally derived from experimental survey results. Where these pertain to Willingness To Pay (WTP) for carbon mitigation, they are discussed in Section 3.3. However, many of these models are simply designed to either predict or describe the decisions empirically without an underlying utility model and shall be discussed here.

One common methodology for analytical, empirical decision models is to break down the decision into two parts. Firstly, a model to determine when a decision maker is likely to be triggered to consider installing a retrofit. Secondly, a second model is created to evaluate the attributes of the retrofit to determine which particular retrofit will be installed. It has also been found that expert advice is a significant factor in both triggering a retrofit decision, and the type of decision made [115]. Decisions have also been found to be influenced by ownership type [116], and socio-demographic makeup [117].

While these decision models are generally based on hypothetical responses by survey participants, Michelsen and Madlener (2014) [117] surveyed approximately 3000 randomly selected homeowners who received BAFA grant for residential heating systems adoption; respondents had been pre-selected for already making a real energy technology adoption decision. In doing so they were able to isolate the drivers for adoption in both newly built and existing dwellings. While new construction energy decisions were more likely to be driven by environmental considerations, such as dependence on fossil fuels, existing building retrofits were driven by the physical properties of the dwelling or the socio-demographic attributes of the decision maker.

There is relatively little retrofit decision model literature relating specifically to the UK.

However, Adan and Fuerst (2015) [118] present a decision model for UK retrofit adoption rooted in microeconomic theory. Instead of using a decision experiment, the researchers use house price data and regionally aggregated retrofit uptake data to build an empirical model of retrofit adoption. The model is based on microeconomic principles and places significant weight on economic factors such as ROI. However, the authors note that the aggregate data on investment decisions is a significant limiting factor in this form of empirical modelling, leaving much of the model construction to future work.

3.2.2.d Alternative Decision Models

While considering retrofits using investment modelling is discussed in some of the scalarised optimisation approaches, it is often captured in values such as NPV or lifecycle cost. This is taken further by Ashuri et al. (2011) [119], who use a sophisticated investment analysis framework as a proposed retrofit decision model. The framework combines an energy model with other sub-models to capture retail energy prices, as well as the investment and regulatory environments that may influence investment decisions. This sophisticated modelling is likely to be more applicable to non-domestic decision makers, however, given the data requirements for evaluating retrofit decisions this way.

Guerrero et al. (2019) [120] present a fuel source transition model using an agent-based social simulation approach. They use a form of bounded agent rationality, including social influence, to model the installation of insulation and a household's transition from natural gas to gas-free heating methods.

Liang et al. (2016) [121] investigate retrofit decisions using a game theory approach, allowing them to analyse the differences between the incentives faced by owners and occupiers. Their analysis determined that reluctance in retrofit installations is often related to the uncertainty of outcomes and the split incentives faced between owners and occupiers of a given property.

While the scope of retrofit decision modelling has been kept to the individual level, there are some cases where domestic dwelling retrofit decisions are made on mass in the case of social housing. Decision modelling of social housing is quite different, with factors such as economies of scale changing the nature of the problem when compared with individual retrofits. While these decisions are outside the scope of this work, Swan et al. (2013) [122] provide a good example of how such decisions are made.

3.2.3 Building Stock Retrofit Adoption Models

HSEMs model the energy use of housing stocks and retrofit decision models can be used to evaluate the likelihood and propensity of a given household to install an energy retrofit; When these are combined they can form housing stock retrofit adoption models (HSRAMs) capable of modelling how the fabric of the buildings in a given stock may change, under certain conditions.

Wang et al. (2018) [112] present a bottom-up retrofit adoption tool designed to evaluate the adoption potential of energy retrofits in Switzerland. Their method, named CESAR (Combined Energy Simulation And Retrofitting), relied on individual EnergyPlus simulations for the HSEM component. This is computationally expensive compared with alternatives but grants higher accuracy, level of detail and control compared with surrogate methods. This energy demand model was paired with a retrofit model which assigned a retrofit statistically, based on an externally provided retrofitting scenario. In this sense, the decision model itself is exogenous, with the uptake scenario determined by the modeller based on historical trends or target uptake levels, rather than on the properties of the decision maker themselves. This study is useful for evaluating the effect of different uptake scenarios but does require the use of an external model, or expert domain knowledge to generate the uptake scenarios. The scale of the analysis is also limited by the computational cost in simulation requirement, with the largest case study presented limited to a neighbourhood of 227 buildings.

There have been some agent-based retrofit adoption models. The most detailed ABM of this type found in the literature is presented by Nageli et al. (2020) [123], who introduce an ABM for energy retrofits including a bounded rational decision model. They pair their decision model with a simulation based HSEM to evaluate the impact of retrofit decisions. Agent decisions are initially triggered by events such as existing property features becoming obsolete. The choice of retrofit installations is based on the investment and maintenance cost of the retrofits, the cost of energy and the agent's WTP. They used representative building agents based on archetypes to achieve scalability, reducing the capacity for heterogeneity and high spatial resolution. The model also foregoes a full optimisation procedure, instead of relying on a small set of potential retrofit solutions selected by the modellers.

The highest resolution retrofit adoption model found was presented by Yang et al. (2022) [124], who was capable of modelling individual buildings using GIS data. However, the resolution of the retrofit installations themselves was very low, with only two scenarios considered. The scenarios were presented only in terms of the performance they achieved, not by the measures that were required to meet those standards. While the model is sufficient to determine the impact of the 'conventional' or 'net zero' standards targeted, no optimisation was performed to ascertain the best performing set of retrofit measures at the level of individual dwellings.

To conclude, existing HSEMs can broadly be broken into top-down and bottom-up methodologies. While top-down models give a good view of general trends and outputs of the entire system, they suffer from low resolution (the inability to investigate sub-components) and can be rigid due to their reliance on historical data. Bottom-up methodologies, such as building energy simulations, can provide significantly more granularity and control to researchers, as well as permitting flexibility such as integration of behavioural or technological components that are not possible using top-down approaches. However, the technical feasibility of modelling entire building stocks using this methodology diminishes with larger building stocks. This is particularly problematic when attempting to perform an optimisation on the stock, in contrast to simply measuring the energy use, as the optimisation procedure may take many iterations. In all the research discovered to date, modelling has required a compromise through either a sampling and upscaling approach based on archetypes or a very limited range of WHRSs that are applied to the entire stock, rather than performing an optimisation [125, 126].

3.3 Carbon Offset and Willingness to Pay

In the field of environmental economics, it is common to model the environment as a public good [127]. The mechanism through which socially responsible consumers interact with public goods, voluntary contributions, are introduced well by Brekke et al. (2003) [128]. Background information on public goods is provided in Section 2.4.2. While early experimental data on public good contributions showed that, particularly in large groups, contributions were low [129], there is a range of circumstances in which contributions are frequent, such as under social pressure or with institutional intervention [130, 131].

We can consider these contributions towards emissions mitigation in terms of the Willingness To Pay (WTP) for a ton of emissions mitigated. WTP is a core concept in microeconomics, behavioural economics, welfare economics, and environmental economics. WTP attempts to capture the maximum amount that a person would pay for something and can be used to allow comparison or evaluation of unobtainable or conceptual goods [132]. This makes the concept useful in environmental economics, as researchers can quantify the value that, for example, unpolluted air might have to a community based on their WTP to prevent the pollution [133]. This is an attempt to quantify the value for something that is not traditionally traded, either because it is infeasible to create a market for, or because, like carbon emissions, it is the byproduct of other market activity.

We focus on WTP as a non-market valuation measure of environmental goods/services due to its applicability to individual household utility functions. It is also relatively achievable measurably and, therefore, has the capacity to be integrated into retrofit decision models. There are, however, a range of other non-market valuation methods which may be used by governmental cost-benefit analyses. As well as consumer WTP methods, these include hedonic pricing models, and travel cost models for public good consumption [134]. The applicability of these methods varies depending on the environmental consideration; house price valuations are often used in hedonic modelling, which can capture non-market amenities such as countryside views. Travel cost models are often used for valuations of public amenities like fishing lakes, as the value placed on these otherwise free services can be calculated based on the costs consumers are willing to incur to visit them. An overview of these non-market valuation methods and best practices can be found in Champ et al. (2003) [134], Bishop and Boyle (2018) [135], and Bateman and Kling (2020) [136] respectively.

3.3.1 Measurement of Willingness to Pay

Measurement methodologies of Willingness to Pay (WTP) can be broadly divided based on two key characteristics [137]. Firstly, methods differ on whether they measure WTP directly or indirectly. Direct methods ask participants transparently what the maximum they would pay for a given quantity of a good or service would be; while indirect methods generally offer a choice of alternatives to the product, whereby the researchers can derive the WTP from the observed decision. Secondly, methods differ on whether they collect actual or hypothetical WTP. Actual WTP involves real choices that affect the participant, either by making them actually pay for a given product or by directly aligning the incentives of the experiment with decisions in the real world. In contrast, hypothetical methods simply pose questions to participants without any real consequences for the choices made [138].

There is significant discussion in the literature about the implications of experimental design choices made when using WTP. Hypothetical design choices benefit from easier survey design and implementation, allowing for questions to be included in larger surveys such as the World Values Survey [139]. The large sample size, international scope, and larger question set in this design allow for statistical analysis that would be infeasible with the practical limitations of incentive-aligned designs [133]. However, some evidence suggests the consumers are more price-sensitive in both real purchasing and incentive-aligned experimental designs than when answering hypothetical questions [137]. This suggests that respondents may indicate higher WTP values when surveyed than in real scenarios where they face similar choices. Nonetheless, some research has indicated little or no measurable difference between hypothetical and actual survey responses [140].

When considering direct versus indirect WTP measurement, there are some further considerations. While direct measurement, using an open-ended question methodology, benefits both researchers and participants through simple and easy to understand design, there are some drawbacks [136]. Direct measurement risks several cognitive biases that appear when participants are asked to evaluate numerically. These include the anchoring effect, in which participants are prone to 'anchor' close to other numerical values to which they were recently exposed [141]. These biases are particularly likely to appear when facing unfamiliar topics [142]. This is likely to make direct numerical assessment an ineffective tool for deriving WTP estimations of WHRSs emissions offset.

A common indirect method of deriving WTP estimations is discrete choice experiments in which participants are offered a range of alternate proposals with different characteristics from which researchers can later estimate the significance of each factor. Statistical techniques such as linear regression can be used to determine the impact of each attribute [143], as well as mitigating some of the methodological drawbacks of direct choice approaches. A good summary of the advantages and best practices of performing discrete choice experiments can be found in Hauber et al. (2016) [144].

3.3.2 Willingness to Pay for Carbon Offset

The use of WTP is suitable for specific environmental amenities such as fishing lakes, or for conceptual components such as a given level of air quality, which is desirable for these localised policies with a specific environmental goal in mind. However, when considering greenhouse gas emissions, the method has been applied to the causal factor directly, measuring the value placed on offsetting or mitigating a given quantity of carbon emissions.

There have been some attempts to calculate the WTP per ton of carbon offset in different settings. Before exploring the most related works in this area, there is some mixed terminology used in the literature that require clarification, specifically the terms mitigation, offset, and capture. Carbon mitigation will, for the purposes of this work, refer to methods, such as retrofits of existing buildings, which prevent future emissions that would otherwise, in lieu of the measures, have occurred. Carbon capture is a term for the capture and eventual storage of carbon emissions from industrial sources, making it a specific form of carbon mitigation. Carbon offset is a more generic term and refers to both mitigation and carbon capture, as well as techniques such as reforestation and afforestation, which are designed to actively reduce environmental carbon, instead of just the emission of it.

MacKerron et al. (2009) [145] used a discrete choice experiment to ascertain WTP for emissions offset of passenger airline flights. The derived WTP per tCO_2e ranged from £10.16 to £38.35 with a mean of £24.26. The most significant contributing factor in offset decisions was certification of the offset, which nearly doubled the derived WTP. Values were negatively affected by the price of the offset and were lower for male than female participants. It is worth noting that the results were only gathered for positive market participants, so this average excludes the proportion of the population who would be unwilling to offset carbon at all. The authors noted that "some caution should be exercised in generalising from the results" [145, p. 1376], as they believed that the sensitivity of WTP derivations may be such that the results may not be robust enough to survive even a change in flight duration, let alone a shift in domain from air travel to the built environment.

Rotaris et al. (2020) [146] performed a similar discrete choice experiment to determine the WTP of 1,228 Italian airline passengers, determining a similar carbon value of between €12 and €38 per tCO_2e. Streimikiene et al. (2019) [147] perform a review of WTP studies and comment on the wide range of factors that influence WTP, including income, gender, and country of origin, while also noting that "pro-environmental lifestyles are related to higher levels of WTP for energy conservation equipment" [147, p. 1480].

While WTP values derived from airline markets have been more widely studied, the universality of the findings may be questioned. Historically, there has been a dearth of literature relating to WTP for carbon offset in the domain of energy retrofits. Some decision model survey designs, discussed in Section 3.2.2, omitted the measurement of WTP to instead focus on non-price factors. Research, such as that of Achtnicht and Madlener (2014) [115] for example, focus on investment triggers and demographic indicators instead of WTP per tCO_2e. While there are few direct WTP derivations found in retrofit decision experiments in the literature, some decision models have discussed or framed their results in these terms.

Banfi et al. (2006) [148] performed a hypothetical discrete choice experiment among residential tenants and property owners in Switzerland. They investigated the demand and WTP for several energy retrofit options, including window and wall insulation at varying levels. No information was provided to participants in terms of tCO_2e saved, an unreasonable expectation as performance models for the participants' dwellings were unavailable. Therefore, while WTP was derived for each retrofit option, these were independent of the qualitative energy performance of the choices. Participants had the greatest hypothetical WTP for new windows, amounting to 13% of their property value compared with the lowest WTP of just 1% to improve standard insulated windows to enhanced insulation.

Adan and Fuerst (2015) [118] present a microeconomic model that incorporates WTP for carbon offset. They consider the WTP for a given retrofit in terms of the financial returns, alongside a weighted coefficient accounting for the 'social cost' of energy use. This value is presented as quite nebulous, encapsulating any non-internalised externalities associated with energy usage. Their conceptualisation is based on a rational economic agent and as such these costs require internalising with a Pigovian tax on emissions [118, p.8].

When it comes to contributions in practice, a 2017 industry report on voluntary carbon offset markets shows an average carbon price of $\$3/tCO_2e$, although this varied significantly depending on the type of mitigation project and the project location [149]. The annual quantity of carbon offsetting is also highly volatile, an indication of an immature market. It is worth noting that these real-world carbon offset prices are significantly lower than reported WTP from consumer studies, possibly indicating a gap between reported and actual decisions. This could also be explained by the high variance in offset prices, indicating consumer willingness to mitigate is much higher than found in industry settings.

3.3.3 Use of Willingness to Pay in Agent-Based Models

Behavioural theories such as those stated above can be found in various ABMs. For example, Andrews et al. (2011) [150] create a utility function for building users that accounts for the warm glow utility received when environmental choices are made. This inclusion of altruistic preferences fell short of fully implemented warm glow theory by the deliberate choice of the modellers, who noted the difficulty of measuring the parameters required to do so. Instead, they used a linear weighting on the disutility caused by cost to others of environmentally harmful action, representing a 'purer' form of altruism than warm glow, and a more traditional approach. Silvia et al. (2016) [151] acknowledge the notion of 'warm glow' environmentalism in their discussion of agent design, but decide against a utility-based approach to altruism, using instead a flowchart decision process that simply assesses the abstract 'environmentalism' of a consumer. The use of social proofing is also common in agent-based models. For example, Snape (2016) [152] modelled the adoption of heat pump installations in the UK using a social model whereby the number of neighbours who adopted the technology influenced adoption rate, although his reasoning was based on domain-specific surveys of adoption rather than based in behavioural science. In their study, Snape acknowledges the apparent lack of behavioural models during a review of agent learning, saying that "[n]otably absent from this discussion of learning algorithms are ... Models based on psychological models beyond pure reinforcement" [153, p.74]. This shows an acknowledgement of a gap for a development framework that focuses on this principle. Rai & Henry (2016) [154] present an excellent overview of how ABMs can be designed for considering consumer energy choices, even recommending the need for chaining formal behavioural rules from the literature to govern agent behaviour. While this review of the ABM process is of high quality, it falls short of being a procedural development process for integrating behavioural components into ABMs, leaving the potential for a more formal framework.

3.4 The Research Gap

Given the interdisciplinary nature of this research, the gap in the literature lies in the integration and application of ML, optimisation, and agent modelling techniques to the physical engineering models used in the existing retrofit evaluation and adoption tools. Section 3.1 identifies the approaches used to model the energy performance and retrofit potential of urban scale building stocks. While there are a variety of techniques for evaluation of the heating demand for these building stocks, they are generally limited in their ability to model the adoption of energy retrofits across the stock. Two particular limits of these modelling methods have been highlighted:

One limitation of existing methods is that the computational cost of bottom-up modelling techniques is, at the scale of an entire building stock, too large to be feasible. This is due to the requirement to perform a large number of optimisation sub-problems. Some building stock retrofit evaluation methods use ML models to overcome this limitation, referred to as Surrogate Energy Performance Models (SEPMs), which allow rapid evaluation of an individual building's performance. However, this technique is too computationally expensive to be used at the scale of large housing stock. Instead of performing WHRSs optimisation, modellers generally pre-select a small set of WHRSs to be applied to the entire stock to evaluate impact, resulting in models that do not capture the WHRSs that would be installed by following a bespoke assessment of the unique dwelling. The technique of constructing a surrogate model that encompasses both the inner SEPM evaluation and the outer optimisation-based decision model, referred to as Surrogate Optimisation (SO), has the potential to rectify this issue by predicting near-optimal WHRSs at a significantly reduced computational cost. However, the method is still in its infancy. The only existing example to be found has limited integration with retrofit decision models, as explained in Section 3.1.4 [82].

The second limitation of retrofit adoption models lies in the use of decision models. Retrofit decision models, which are analytical models designed to capture the behaviour of retrofit decision makers, have been attempted in the literature. However, there are relatively few as discussed in Section 3.2.2. While some models attempt to calculate the WTP for a given retrofit or feature, as described in Section 3.3, none are presented in terms of transparent measures of environmentalism, such as WTP per ton of greenhouse gas emissions mitigated. There have been some studies in other fields, most notably voluntary carbon offset in the airline industry, which investigate environmental decisions in this framing. This problem formulation would allow the construction of decision models based on voluntary contributions to public goods, as described in background Section 2.4.2. This allows for multi-objective optimisation-based modelling of uptake among rational and environmentally conscious agents. This has not been found in the literature, either in traditional adoption methodologies or combined with SO techniques.

To conclude, there is a gap in the literature for a high-speed residential building stock retrofit adoption modelling methodology, based on the SO method, which can be integrated into WTP-based retrofit decision models. This novel combination of ML with agent-based decision models would allow for a higher level of detail in modelling than is currently technically feasible in housing stock energy models for retrofit adoption. A proposal for how to reframe the retrofit decision models in this way and integrate the model into a trained SO is put forward in Chapter 4, then implemented and explored in later chapters.

Chapter 4

Methodology

It was established in Section 3.4 that there is a gap in the literature for a method of computationally feasible retrofit adoption modelling with integrated environmentally conscious agent decision modelling. In order to fill this gap, a machine learning approach will be taken, based on integrating a Surrogate Energy Performance Model (SEPM), an optimisation process and different models of agent utility. This method, which is referred to in this work as Surrogate Optimisation (SO), will then be integrated into an agent based retrofit adoption model. These sub-components, and their scopes, are shown in Figure 4.1. As the contribution is methodological in nature, specific implementation

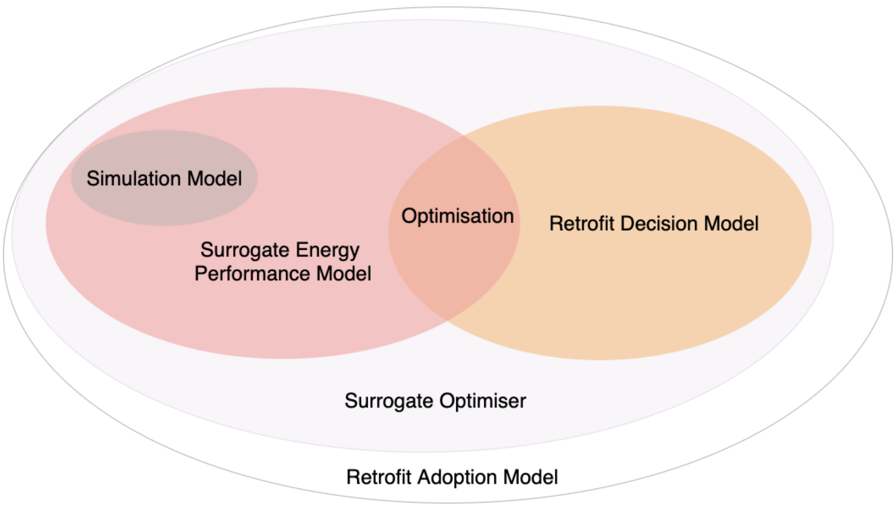

Figure 4.1: Venn diagram showing the interaction of scope between the sub-components of the final retrofit adoption model method.

details of the research will be introduced alongside the results presented in later chapters. The purpose of this chapter is, therefore, to provide a high level description of the stages of SO, as well as to introduce the data and software tools to be used throughout the research.

Section 4.1 describes the methodological approach to the agents used in this research, including a high level description of how the decision models used will be formed, from a theoretical perspective. Section 4.2 describes the methodological steps involved in the construction and training of the Surrogate Energy Performance Models (SEPMs) used in this research. Section 4.3 follows on from this, discussing how the SEPMs will be integrated with optimisation processes to obtain near-optimal retrofit solutions for individual buildings. Section 4.4 lays out the approach for training a Surrogate Optimiser, based on a data set of values generated using the SEPM based optimiser combined with the different decision models of interest. Section 4.5 will introduce and investigate the primary building data set used throughout the research. Finally, Section 4.6 will describe the software tools used, including the EnergyPlus based simulation framework used for the simulation based energy evaluation stages.

4.1 Rationality of Agents

One key component of building stock retrofit adoption modelling is the retrofit decision model. The ways that researchers have approached these models is discussed in Section 3.2.2, with approaches ranging broadly depending on the intended purpose, scope and level of detail of the research conducted. Through this work, the decision model will be framed in terms of agent decisions and, therefore, the models that contain them will be referred to as agent-based models, the background for which are discussed in Section 2.1.4. Generally, the agents used will be a form of utility-based agent, with an objective function made up of different attributes that contribute to their decisions. Different assumptions about behaviour will be implemented, justified and at times relaxed to increase the complexity of decisions.

4.1.1 Purely Rational Agents

The notion of rationality and rational choice is quite broad but we will use a broad definition of an agent taking actions that, to the best of their knowledge, maximise a given goal that they have [155]. This places rational agents in agreement with both goal-based and utility agents discussed in the agent-based simulation field in Section 2.1.4. However, this broad definition of rationality can be narrowed to a classical microeconomic consideration of a self-interested agent, which we will refer to as purely rational. The purely rational agent is only willing to consider the benefit they themselves receive from a given decision.

$$NPV = C_0 - \sum_{t=0}^{n} \frac{R_t}{(1+i)^t} \tag{4.1}$$

Equation 4.1: Net Present Value for an investment of C_0, returns of R at time t and a discount rate of i.

The advantage of a rational framing that combines rational choice and utility theory is the ability to scalarise the problem into a single-objective optimisation problem with a utility function, as discussed in Section 2.3.2. To construct an objective function for a purely rational agent we take the Net Present Value (NPV) of the energy bill savings made over the period of the investment, less the cost of the retrofit. A background discussion of NPV can be found in Section 2.4.3 and the formula is presented in Equation 4.1. NPV is established as a viable decision tool in the investment and adoption literature [54, 55, 56]. It has been criticised for not accounting well for future changes in technology and thus failing to account for some possible opportunity costs that can be captured by waiting [57].

One of the key features of the NPV is the discount rate. Conceptually, the discount rate attempts to capture the opportunity costs and inflation that occur between now and future income. While there are different methods of discounting future value, the general standard is a flat compounding discount rate, signified by i in Equation 4.1. In this research, an approximation of the UK's risk-free rate of return is used for the discount rate, which will be discussed during implementation.

NPV was selected as the objective function for purely rational agents for several reasons. As discussed in Section 3.2.2, NPV or some close equivalent, such as lifecycle costs, make frequent appearances in retrofit analysis tools. These are particularly common metrics when considering decision frameworks or more institutional decision makers due to the measurable and clear financial nature of the objective. Another advantage of NPV is that it captures both the upfront financial cost of the investment, the ongoing costs, and the ongoing benefits in a single value, which can be expressed in present value currency terms. This allows for a very simple decision model, as a rational agent would take any free choice with a positive NPV (if the alternative choice is nothing).

There are alternative scalarisation techniques that could be considered, instead of NPV. Sometimes LifeCycle Cost (LCC) is considered in the literature, although typically this is used as a secondary objective when accounting for other financial or comfort considerations. Some references to LCC are actually net LCC and factor energy cost savings into the modelling, making them functionally equivalent to NPV. An alternative, perhaps more economically sound but significantly less feasible performance metric may be expected utility. In contrast to NPV, which just discounts future values at a flat, compounding rate, expected utility attempts to also capture the difference in utility that a given cost or saving will have due to other factors in an individual's life. Expected utility theory would include, for example, an income effect, whereby savings are more beneficial, and costs more harmful, to an individual with a lower income. This is due to the core economic concept of diminishing marginal utility. These factors will be considered outside of scope as they add a level of complexity, and data requirement, about decision makers which is not just unavailable but possibly immeasurable. However, with some assumptions about the agents, which will later be relaxed, NPV based investment decisions can be aligned with expected utility theory.

In order to consider NPV to be the decision metric of purely rational agents, some assumptions must be made. A list of these initial assumptions can be found below, followed by a discussion or justification. Further assumptions, scope limitations, and conceptual modelling will be introduced elsewhere. When discussing the strength of an assumption, strong assumptions can be considered those which abstract significantly from reality and require the most justification. In contrast, weak assumptions only deviate slightly from the real-world system and are, therefore, less likely to introduce error.

List of assumptions made:

1. Agents are self-interested and ignore environmental impact.
2. The installation can be purchased outright or finance is available at discount rate.
3. Future savings and costs can be predicted or provided to the decision makers before installation.
4. Thermal comfort is achieved through setpoint and scheduling alone.

Assumption 1 has already been discussed: that agents are only self-interested. This will be relaxed when modelling environmentally conscious agents as discussed in Section 4.1.2, and potentially relaxed further when considering agents as uninformed decision makers in Section 4.1.3.

Assumption 2 is only an assumption insofar as we use a positive NPV as an entire decision model. Where we use the NPV objective in conjunction with a trigger model, for example, then this assumption is relaxed. This can therefore be considered an assumption only of the most optimistic, retrofit potential models, as opposed to more grounded retrofit adoption models for which this assumption alone would be too strong.

Assumption 3 assumes a high degree on information available to decision makers. A decision model based on optimising NPV assumes both that the NPV of any given retrofit can be calculated for a given installation of a given dwelling, but also that the value is being optimised by the decision maker. In practice, this is really modelling two separate processes. The optimisation, or a close approximation of it, would be performed by an expert from a retrofit installer. The near-optimal solution, along with the salient information, would then be presented to the property owner or decision maker when making a quote.

Assumption 4 relates to the significance of thermal comfort to the decision maker. This assumption can be broken down further to first assume that overheating, the propensity of the temperature in a dwelling to become unreasonably warm, is not a primary decision factor. Given the weather profile in the UK, which is the scope of this problem, this assumption is fairly weak. However, as global climates change this could be more significant in the future. The second part of this assumption is that the heating systems in dwellings are sufficient to meet the demand required by the heating setpoint, in a timely manner. To account for this assumption, it will be made clear that estimates are made for heating energy demand, rather than energy use. This accounts for deviations between the energy required for a system to achieve a setpoint, compared with what the system is capable of outputting. This will, however, be controlled for during simulations by capping the heating output capacity to a reasonable level, weakening this assumption further.

4.1.2 Environmentally Conscious Agents

While modelling rational agents may be useful, particularly if we wanted to consider a base case adoption scenario where only financially viable retrofits are performed, it is lacking in several ways. It is observable in both experimental and real-world settings that individuals will act altruistically in an attempt to contribute to the environment. The background of this phenomenon is discussed in terms of public goods games in Section 2.4.2. The prior research relating to the measurement and modelling of agents who are willing to pay for a reduction in carbon emissions was reviewed in Section 3.3.

$$U = f(v^{+}, e^{-}) \tag{4.2}$$

Equation 4.2: Generic utility function, with the NPV (v) of the investment increasing utility, and personal carbon emissions (e) decreasing utility.

Relaxing the assumption, based on evidence that consumers are willing to forego some personal financial gain in order to contribute to emissions reduction, will change the objective function that rational agents are attempting to optimise. Generically, we can consider a utility function, shown in 4.2, which is positively impacted by NPV (v) and negatively impacted by the environmental emissions (e). While this relaxation alone is uncontroversial, it is the nature and formation of the utility function that is highly contested within the literature. This will require us to make some assumptions about agent behaviour in order to model them. As discussed in Section 2.4.2 there are several theories as to the behavioural principles that drive these environmental contributions with which the chosen utility function must be in agreement.

$$LCS = \sum_{t=1}^{t=p} e_t - E_0 \tag{4.3}$$

Equation 4.3: Lifecycle Carbon Saving (LCS) calculation.

Before presenting candidate utility functions for environmentally conscious agents, it may be beneficial to consider two types of decisions that are to be made and reframe our discussion of environmental decisions in light of energy retrofits. In the same way that we used NPV to scalarise the financial impact of a given retrofit investment, we can use Lifecycle Carbon Savings (LCS) to represent the GHG emissions of a given investment. LCS is represented by Equation 4.3. The value of emissions saving, e_t, will depend mostly upon the energy saving properties of the retrofit. It will also be impacted by fuel type, as the emissions for gas and electric energy generation differ. The lifetime embodied carbon is represented by E_0 and considers the cost of manufacturing, transporting and disposing of building materials used in the retrofit. LCS can be considered the GHG emissions equivalent of NPV, with the key difference being the lack of discounting. The use of carbon discounting was considered to weigh near emissions more heavily than future emissions but was not implemented as it was not a common methodology found in similar research.

$$U = f(v^{+}, e^{-}) = v + \gamma s \tag{4.4}$$

Equation 4.4: Linear utility function, with the NPV (v) of the investment increasing utility, and LCS increasing utility by γ per ton of emissions saving. $CO2_e$.

The agent's utility for a retrofit solution is shown in Equation 4.4, with the household's carbon valuation expressed as *gamma*. Given the *gamma* is expressed in £$/tCO_2e$, the utility can be expressed in present value financial terms, allowing a direct trade-off between the terms. This carbon valuation, *gamma*, represents the WTP for carbon mitigation. In a linear model, this is also the coefficient of marginal utility with respect to emissions mitigation.

The driver for this value can be seen as the personal or social components that motivate households to be carbon conscious and could be driven by a combination of altruism, social pressures, or self-preservation. This utility curve formation is generally agnostic to the underlying cause of the contribution, only accounting for the fact that the climate contribution has value, instead of the underlying cause of the value.

The proposed utility function is linear in both financial and carbon consumption. This contradicts a general tenet of many utility functions, diminishing marginal utility. This property gives utility curves a concave shape, which is a useful element in many forms of analysis. This assumption is relaxed for two reasons. When relating to general economic consumption, diminishing marginal utility is caused by the income effect: as income rises, the utility benefit of additional income diminishes. It can be relaxed when considering the NPV from retrofit investments due to the relatively small size of the financial benefit, and due to the long time period over which it is spread. This allows us to use a point approximation of utility, as the size of the income effect is small.

The strength of the assumption of constant marginal utility from carbon emissions reduction depends on the underlying cause for the public good contribution. A model for pure altruism is likely to result in small but linear marginal utility. Socially driven conspicuous altruism may lead to drastically diminishing marginal utility from carbon mitigation (as the appearance of mitigation dominates the volume of emissions reduced). However, the linear framing has some major benefits, mostly in the ability to directly survey the financial equivalence of a given degree of emission saving, something which can be ascertained more readily than a convex curve. This assumption can also be justified similarly to that used for the NPV. Given the long time span of the retrofit, it is likely that emissions offsets are small relative to the total carbon footprint of the household when factoring in travel, food, and transportation of consumer goods.

4.1.3 Uninformed Agents

Another concept pertaining to agent rationality that is important to capture is that of imperfect information. It is reasonable to assume that agents performing a retrofit optimisation are doing so with the help of a professional and are therefore aware of both the availability of retrofit options and the impact that those installations are expected to have on their dwellings. However, it would not be reasonable to assume that all household agents are aware, a priori, that there may exist positively performing WHRSs to be installed. To rectify this, a trigger model is to be used in conjunction with the optimisation procedures to determine when agents will choose to evaluate the retrofit potential of their dwelling. Trigger models are found in various related works, including several HSEMs described in Section 3.2.

The nature of the trigger model used will vary throughout the work. Chapters 5 and 6 will use a homogeneous and endogenously provided probability that a given agent will be triggered over a certain time period. This is a top-down approach to modelling triggers, relying on observation of a real-world aggregate value (the number of retrofit evaluations that occur per year) and applying it across the agent set. The focus of Chapter 7 will be on replacing the trigger model with a more sophisticated bottom-up approach. This will be done using survey data to determine the relative importance of different potential triggers as perceived by a set of household archetypes. The absolute probability of each trigger leading to a retrofit evaluation will then be calibrated using the exogenously provided base rate. While still relying on a base rate for calibration, this will allow for heterogeneous household archetypes as well as a bottom-up design that will allow altering of the occurrence rates of potential triggers to allow for more detailed scenarios to be modelled. The methodology behind this process will be laid out in more detail in Section7.3.1.

4.2 Surrogate Energy Performance Modelling

The methods for evaluation of the energy performance of dwellings, including potential retrofits applied to those properties, are discussed in Section 3.1.1. While there are some analytical models available, the flexibility and accuracy of simulation modelling were

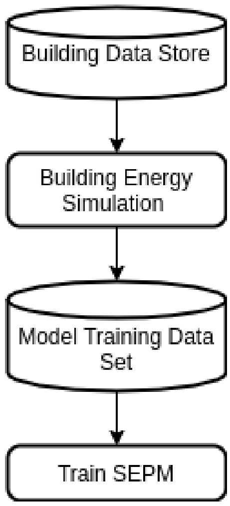

Figure 4.2: Generic and high level overview of the SEPM training process

found to be a common method. A background discussion of building energy simulation was provided in Section 2.3.4. While it was found that simulation modelling carries many advantages over rigid analytical models, it also carries a significantly increased computational cost. This computational cost can be overcome with the use of Surrogate Energy Performance Models (SEPMs), a technique that has been used to leverage the advantages of simulation modelling while mitigating the computation costs. A discussion of related works which use SEPMs can be found in Section 3.1.2.

A generic and high level flow chart showing the basic components of a SEPM is shown in Figure 4.2. The concept behind the surrogate model is quite simple, generating a training set using the simulation software. This training set is used in conjunction with a ML algorithm, typically a regression model, capable of estimating continuous objective values of the type discussed in Section 2.2.2. In practice, each stage of surrogate model formation requires careful consideration such as simulation model construction, selection of modelling technique, model features, and training hyperparameters. These implementation steps will be discussed in more detail in Section 5.2.4 when the primary SEPM is trained and used for energy evaluation.

The primary benefit of using a SEPM, compared with direct energy simulation, is the speed of making predictions using trained models. This allows for increased computational feasibility when performing retrofit optimisations, as discussed in Section 4.3.

An import process in SEPM training is modelling method selection. When evaluating the related literature (Section 3.1.2), there were a variety of modelling techniques used. The most common method was linear based regression, but more recent work, and work that focused on optimisation was focused more towards Artificial Neural Networks (ANNs), including Deep Neural Networks (DNNs) [79]. This is likely due to the flexibility of DNNs in modelling a more varied set of non-linear inputs, such as building archetypes, as input features. However, while the modelling technique is likely to be important, it has been found that a broad set of techniques can be suitable for SEPM applications, dependent on the specific problem but without a specific algorithm dominating the problem domain [76].

An important process in SEPM training is feature selection. The initial feature selection stage involves collecting potentially valuable model features using domain knowledge, as well as insights from SEPMs in the literature. This is likely to include direct properties, such as the heating setpoint, derived features such as floor space, glazing ratios (of different parts of the building), and details of the retrofit status of the building. These features will be discussed in more detail at the implementation stage.

After initial feature selection, backwards elimination of feature selection will be used to remove any poorly performing or unnecessary features. This stage can both increase accuracy and reduce the complexity and data demands of a model. A statistical test will be performed on input features, eliminating the most poorly performing ones until the elimination of any remaining feature reduces model accuracy in a statistically significant way [156]. The details of this test, and the features eliminated, will be identified at the implementation stage.

Another important part of SEPM construction is to determine the target values of the model. While all SEPMs will evaluate the energy performance of the building, there are several target metrics that could be considered. Although most examples in the literature focused on energy [75, 72, 71], there were also models for thermal discomfort and cost [74]. Given the utility functions of decision makers discussed in 4.1, the annual heating energy demand is likely to be a sufficient objective value from the surrogate model, with the cost and carbon emissions savings calculated with post-processing outside of the model.

Given the continuous nature of the annual heating energy demand, expressed in gigajoules (GJ) or kilowatt hours (kWh), selected models will be regression models, the background for which are discussed in Section 2.2.2. When evaluating the performance of these models, continuous evaluation metrics will be presented, including the coefficient of determination (r^2). The r^2 value represents the proportion of the variance in the target value which is explained by the input variables of the model, with a value of one representing a perfect model. This metric is highly relevant in evaluating the model performance [157]. However, r^2 has a flaw, as overspecified models are not punished and so r^2 values can be improved by adding random and uncorrelated variables into the model, a problem which is mitigated by the use of the adjusted r^2 metric, which reduces in value as additional model variables are added [158]. As such, adjusted r^2 values will be presented with regression model results. While r^2 is a useful summary statistic, it is presented as a fraction of the variance of the target value. Metrics such as Root Mean Square Error (RMSE) and Mean Absolute Error (MAE) present the error in model units, giving a better intuitive understanding of the size of the errors introduced by the model. Where possible, the distribution of these errors will also be presented and visualised in the research, demonstrating the model performance across the range of relevant values.

The choice to use a SEPM, as opposed to an expert-derived analytical model (such as BREDEN [60]) was due to the flexibility of simulation based models. As discussed in Section 3.1.1, there is some conceptual equivalence between trained SEPMs and expert derived analytical models. The trained models are analytical in nature and so benefit from fast evaluation speed. Similarly, expert derived analytical models will be highly data driven, and are likely to have been guided by the simulation in their construction. The main advantages of SEPM come from the existence of the simulation framework used to train it and from the flexibility in the derivation of the model. The existing simulation framework allows for the error of the SEPM to be known precisely, compared with the simulated building. This gives a measurable degree of error compared with analytical models which are more opaque. Similarly, the simulation framework gives flexibility to SEPM data requirements, allowing for feature engineering based on the available data, rather than having to work to the data expectations of the original model designers. Even models like RdSAP, which were designed to lower the data requirements compared with the full SAP, expected practitioners to have some level of access to the property to acquire the necessary input data [62].

4.3 Optimisation using Surrogate Energy Performance Models

It is often desirable, when performing building stock models, to perform an optimisation as part of a decision model. This is the case for decisions made by the utility-based agents discussed in Section 4.1, who are attempting to make decisions that optimise their personal utility, as determined by the performance and other properties of a given retrofit. Optimisation of a single building's energy retrofit potential can be a computationally complex task, as discussed in Section 3.1. When attempting to optimise a whole

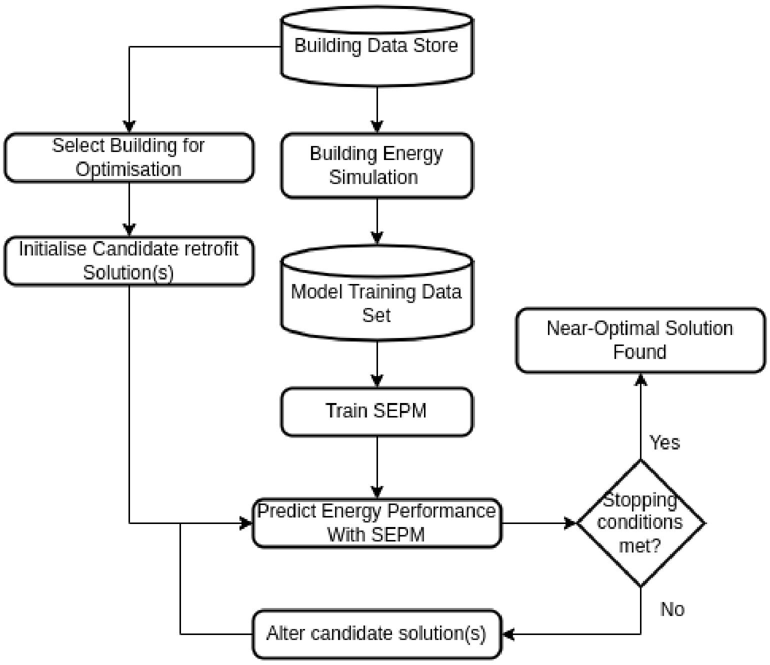

Figure 4.3: Generic and high level overview of Optimisation using meta-heuristic optimisation loops with Surrogate Energy Performance Models.

building stock, the complexity increases at least linearly, making the computational burden of the task significant. The use of SEPMs in optimisation, which is generically described, at a high level, in Figure 4.3 allows for more rapid evaluation of potential solutions, reducing the computational cost of optimisation significantly. As discussed in Section 3.1, this method is still insufficiently fast to allow for large scale evaluation of an entire building stock, with most related works optimising single buildings or small stocks such as a neighbourhood.

The process of optimisation with a SEPM involves two key stages. The first stage was outlined in Section 4.2 and involves the use of simulation software to prepare a data set for the training of the SEPM. The second stage is the embedding of this SEPM as an evaluation method inside a metaheuristic optimisation process. The process is conceptually agnostic to the optimisation metaheuristic chosen, although given SEPM based optimisation is designed to increase evaluation time, there will be greater benefits in optimisation scenarios that require a larger number of function evaluations. A background discussion of optimisation, including heuristic optimisation, can be found in Section 2.3. The choice of optimisation method, its hyperparameters, and performance will be discussed at the implementation stage, which can be found in Section 5.3.

One drawback of using SEPM based optimisation methods is the introduction of error between the simulation model and the SEPM. While this is the case for any form of model, it may be exacerbated in optimisation if the errors are not evenly distributed, resulting in over-representation of solutions in spaces that are more error prone. Consider, for example, a SEPM that reports good accuracy on average, but is prone to overestimate the energy savings of highly insulated window panels. This may not present as a significant problem during SEPM training, as the average error is low and these panels are rare in the data set of existing buildings. However, this could lead to a large number of optimisation solutions containing these panels which, in a simulation based optimisation would not be common. This highlights the importance of examining not just the average error of the SEPM but also performing analysis of the distribution of errors and the sensitivity of the model to different input changes.

4.4 Surrogate Optimisation

The process of Surrogate Optimisation (SO), using a data set of near-optimal retrofit solutions to create a predictor of retrofit solutions, is laid out generically in Figure 4.4. The limited examples of this method found in the literature are laid out in Section 3.1.4. This process is referred to in this work as surrogate optimisation due to the similarity with the surrogate modelling technique which is expanded upon as part of its construction. This can be seen by the entire subsumption of Figure 4.3 in the SO method overview, as the SEPM based optimisation method is used to form the data set for SO training.

The model selection process here is likely to be dependent on the nature of the retrofit solutions estimated. This research will consider both discrete and continuous formulations of retrofit components, as both the materials used and the thickness of said materials are to be selected by retrofit decision makers. As such, model selection will require the ability to perform regression and classification predictions. This can be done by using separate modelling techniques or by post-processing predictions. Binning can coerce continuous values into discrete ones, while discrete values can already be expressed continuously.

The feature's available for the SO model will be similar to those available in the SEPM, given that the SEPM is used in the energy analysis performed in the prior optimisation stage. As was shown in Figure 4.1, However, the SO model also encapsulates the retrofit decision model, and therefore the assumptions of agent rationality discussed in Section 4.1. In the case of rational agents, whose objective is to maximise the NPV of the investment decision, features of the investment scenario are, at the point of SO training, either locked into the model or to be included as features. In the case of envi-

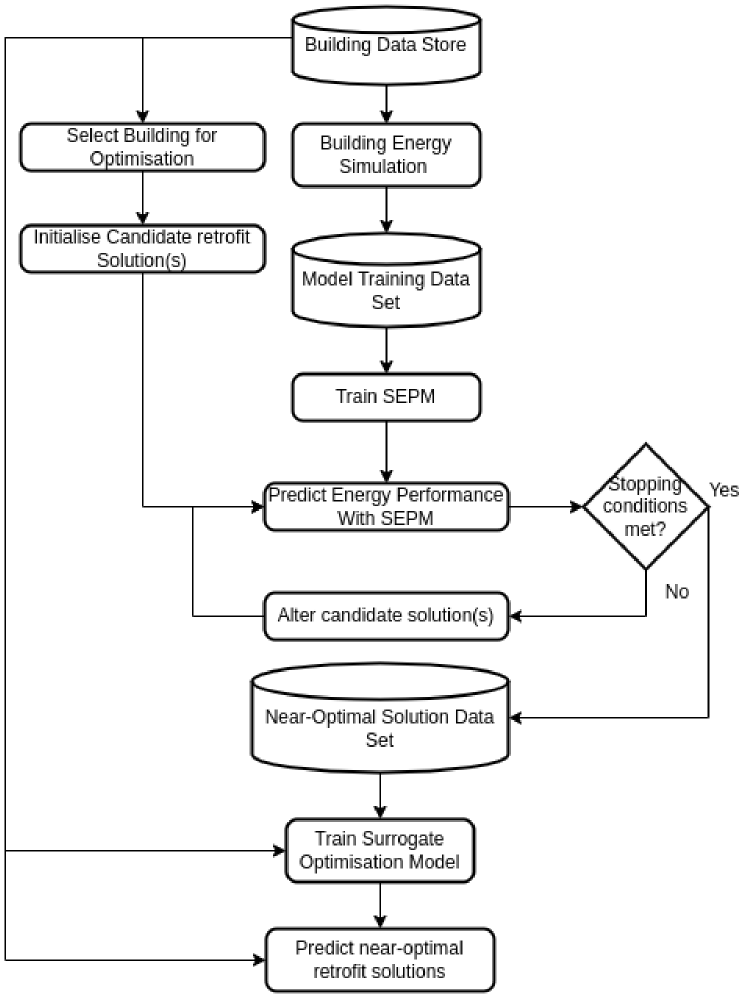

Figure 4.4: Generic and high level overview of Surrogate Optimiser training process.

ronmentally conscious agents, model features will also include properties of the retrofit decision agents and will include their marginal WTP for carbon reduction, as discussed in Section 4.1.2.

When considering the purpose of the SO method, it is worth considering whether the approach laid out thus far is sensible. Given we are dealing with agent decisions, a reinforcement learning approach of the type considered in Section 2.2 would be more appropriate. However, reinforcement learning focuses on sequential and bidirectional interactions with a given environment. The process itself is not a replacement for optimisation techniques of the type discussed in Section 4.3 and would result in training agents with a specific goal based on their characteristics, rather than a generalised pattern for how an entire population of agents would make decisions based on the rules provided.

A methodological comparison can be drawn between SO and model based methods for optimisation, discussed in Section 3.1.5. Both use similar tools: constructing a model of an underlying solution space for the purpose of predicting near-optimal solution, there are some differences. Model based methods are used as optimisation metaheuristics of the type discussed in Section 2.3, used to generate approximations of a specific optimisation problem to expedite the search for a near-optimal solution. In contrast, SO uses a data set of a priori near-optimal solutions alongside features of the decision model the optimisation depends upon, to predict solutions that real-world decision makers would use. This allows for the rapid evaluation of the behaviour of heterogeneous agents which would otherwise require individual optimisation procedures.

4.5 Building Data Set

One of the primary data sets used in this research pertained to 95,500 dwellings in the city of Nottingham, UK. A GIS visualisation of the data set can be seen in Figure 4.5, demonstrating the scale of the data, the location of buildings in physical space and the archetype labels applied. The complete set of features in the building data set, with a brief description of each, can be found in Section A.1 of the Appendix.

Some of the pre-processing stages required to create the database were performed prior to this research by members of the research team, the details of which can be found in Julian et al. (2019) [103]. These included archetype generation, dis-aggregation and attribution of properties from the English Housing Survey (EHS) using stochastic assignment from a cumulative density function. These were combined with processing of physical attributes from OS map data.

4.5.1 Total area

The total area value, calculated by multiplying the footprint area by the number of stories, will be a key summary statistic, particularly in the construction of the surrogate energy performance model later. It is also a value with the most outliers detected in Section 4.5.3, so we will examine the distribution of values in greater detail here. The distribution of the data, which can be seen most clearly in Figure 4.7, is bi-modal with

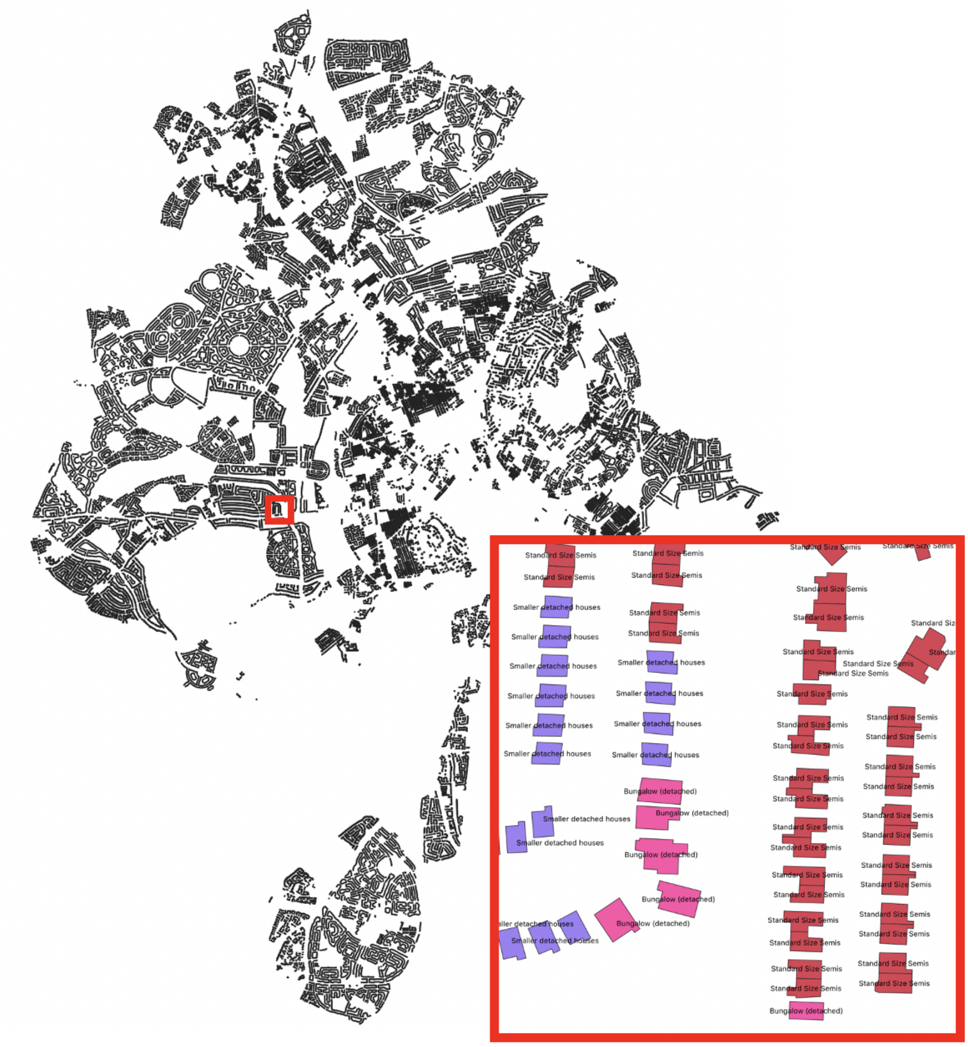

Figure 4.5: GIS visualisation of the building data set, with a small portion highlighted to show archetype allocation.

a long right tail, indicating two main dwelling size clusters as well as a small number of very large dwellings. These large dwellings are likely to fall into either errors (buildings that are incorrectly calculated due to imperfect labelling) or buildings out of scope such as blocks of flats or large historical buildings. These will be discussed more in Section 4.5.2. These anomalous or out of scope values account for some of the right tail, although some can be accounted for by the limits on the left tail, since there is a hard lower limit on properly size (which cannot be negative) but also a solid practical limit to what can constitute a residential dwelling. The bi-modal data is also largely explained when breaking the data down into the building archetypes which compose it, which are generally uni-modal as shown in Figure 4.6. This also shows the importance of the archetype distinctions, as when construction models of the heterogeneous population of building types, we must ensure our models are able to account for these different underlying distributions.

4.5.2 Data Cleaning

Some initial data cleaning was performed on the data set described up to this point. There were broadly 3 main drivers to each cleaning operation that was carried out:

- Error correction
- Scope management
- Data transformations

When using data allocated from a CDF, the notion of error is somewhat interesting. As values are statistically assigned, we may only expect them to be correct on average, with errors cancelling but not actually incorrect. These assigned variables are also not susceptible to outliers in the same sense that collected data may be, as any highly divergent values are drawn from a defined CDF. However, some values in the data set, particularly relating to building features or archetype categorisation, are more likely to be errors that require removal. The most common errors found were due to buildings that were categorised in a way that severely altered their features, such as assigning a single building TOID to an entire row of terraced housing, giving it an infeasible size and shape for a single dwelling. The inverse was also noted, with the accidental assignment of residential building categories to TOID objects like bus shelters, which are too small to conceivably be residential properties. These errors, lying at the extremes are easier to identify with the methods discussed in Section 4.5.3.

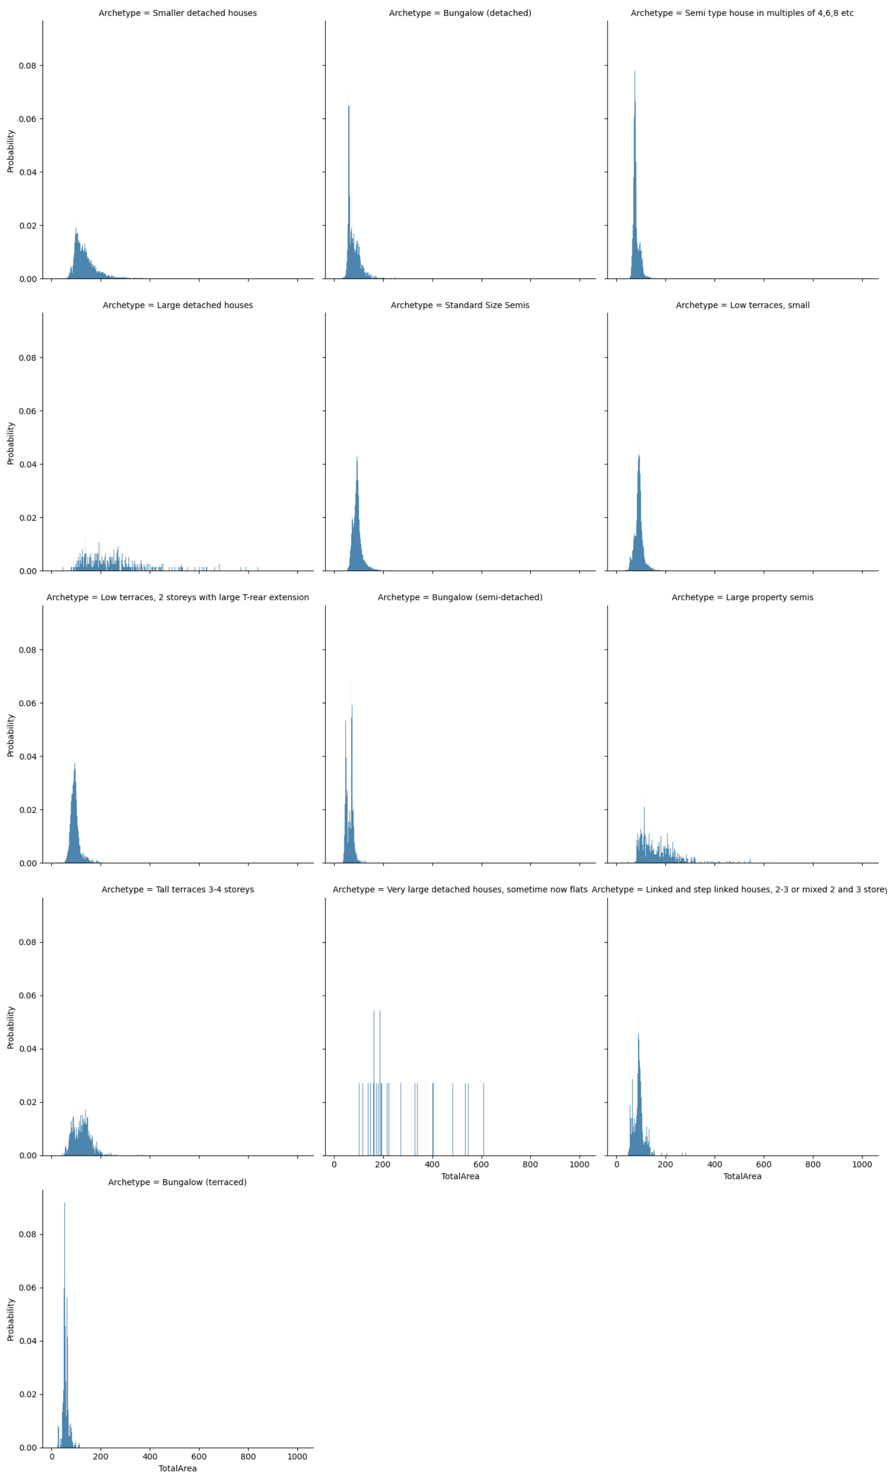

Figure 4.6: Histograms of Total Area stratified by building archetype.

Some errors, which do not result in anomalous features, are significantly more difficult to identify. For example, incorrectly labelling a terraced house as semi-detached. However, these errors are likely to be less significant.

When it comes to scope management, there were more decisions to be made at the data cleaning stage. While the scope had already been limited to residential properties only, it was also decided to remove any flats and mixed-use dwellings from the data set for the entire course of the research. The reasons for this decision fall into practical and theoretical categories. From a practical perspective, there is no data available for internal partitions within the blocks of flats which appear in the data set. This cannot be accurately estimated from the available footprint data. The errors arising from this uncertainty would compound given the higher level of thermal interaction between flats compared with other dwelling types, due to an increased number of shared surface interactions. When also considering the prominence of unheated communal areas in this dwelling type, the level of uncertainty in the data was deemed too high for the generalised model attempted. An additional, more theoretically driven reason for omitting this dwelling type was a hypothesis that the decision making process was likely to be different for this property type, with a larger number of short-term tenants as well as single-owners in large blocks, the retrofit decision landscape was likely to be very different and with a more corporate mindset focused with a greater focus on economies of scale. While a worthwhile area of study, it exists outside of the household level scope maintained throughout this work.

The data transformations relate to the pre-processing techniques that improve model performance. The importance of these techniques depends on the model selected, and the specific transformation or set of transformations used can be discovered during the model tuning stage of a model creation process. As such, these will be discussed throughout the thesis when relevant. They include techniques such as min-max normalisation, standardisation, principle component analysis, and feature selection.

4.5.3 Standard score anomaly detection

$$z_i = \frac{x_i - \mu}{\sigma} \tag{4.5}$$

Equation 4.5: Standard score formula, where x_i is the value of an observation, μ is the mean and σ is the standard deviation.

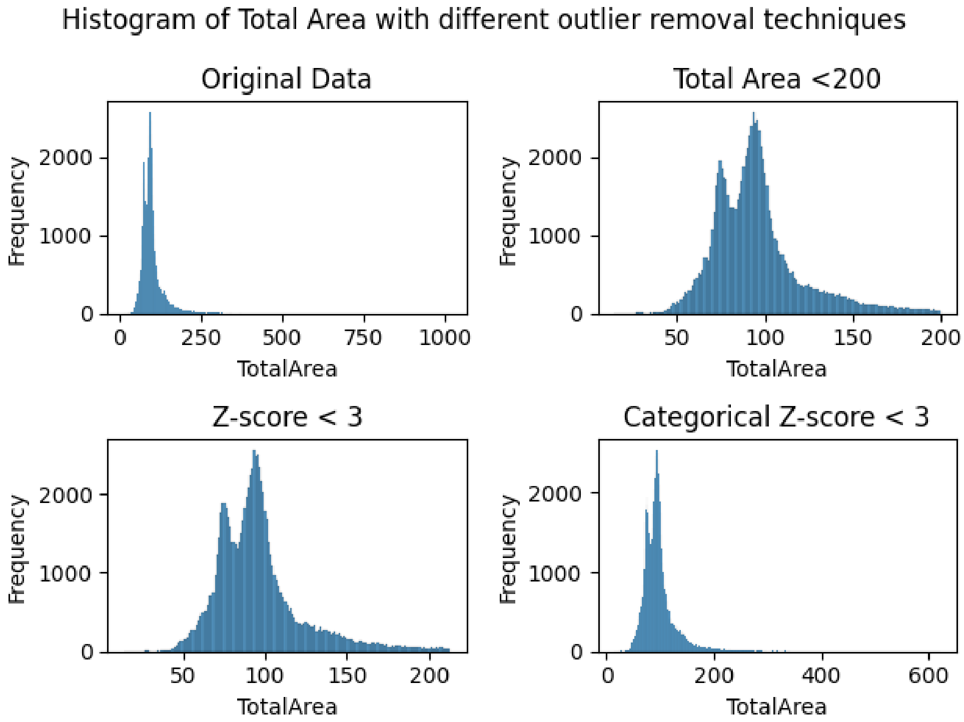

Figure 4.7: Histograms of Total Area with different outlier removal techniques.

The standard score, sometimes referred to as z-score, reports the number of standard deviations that a data point lies from the mean, as shown in Equation 4.5. This is one method of detecting outliers by creating a z-score threshold above points that will be removed as anomalous. There is, however, a flaw in blindly applying this technique to our data as certain building properties are drawn from different distributions depending on the archetype. In our raw data, for example, the mean area of a terraced bungalow is 58 m^2 while large detached houses had a mean floor area of 246 m^2. Both of these archetype means differ significantly from the overall population mean of 99 m^2. In fact, if you apply a z-score elimination of outliers to the entire population, the upper threshold value (with $z <= 3$) would be 213 m^2, below the median value of large detached houses. This would result in over half of this archetype being removed from the data set. As such, we have evaluated data points based on a categorical z-score: adjusted for the attributes of each building's archetype. To see the impact of this adjustment, see Figure 4.7 which compares histograms of total area with different elimination strategies.

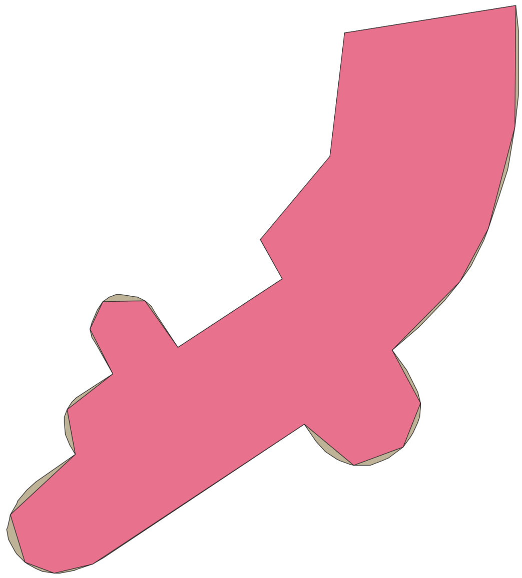

Figure 4.8: Results of polygon simplification algorithm on building osgb1000022745085.

4.5.4 Polygon simplification

A process of simplification was performed on the polygons at the pre-processing stage. Given that buildings are defined in 2 dimensional polygons, certain architectural detailing, such as curved walls, are represented in the original data set with a large number of small polygons to create the curve. This is inefficient for simulation, data storage, and IDF processing. These tiny polygons also cause issues for the placement of glazed surfaces and any derived values such as average wall size.

The GRASS vector module v.generalize [159] was used to simplify shapefiles, using the Douglas-Peucker algorithm [160] with a tolerance value of 0.5m. This parameter value was chosen to result in minimal changes whilst dealing with the more unusual buildings, resulting in a relatively small number of polygon changes (4.1%). The effect of this simplification can be seen in Figure 4.8, which shows an example of a larger and more complex dwelling that is simplified without changing the fundamental structure. In this case, the footprint polygon was reduced from 117 to 26 edges. This is an extreme example of simplification, with a mean edge reduction of 9.2 in the subset of simplified buildings.

4.6 Toolsets

4.6.1 EnergyPlus Framework

As discussed in Section 2.3.4, the energy simulation software used in this work is EnergyPlus. In order to run an EnergyPlus simulation, the software must be provided with three files. The EnergyPlus Weather (EPW) file, the Input Data Dictionary (IDD), and the Input Data File (IDF).

The first data file required for EnergyPlus is the IDD, which defines the components that can be used in the software. This file is used in the development and extension of EnergyPlus and was not altered in the process of this research.

The EPW file contains the necessary meteorological data to simulate the energy performance of the building. The data was sourced from the Met Office weather centre located in Watnall (station id 03354), approximately 10 kilometres from the centre of Nottingham [161] and generated using the UKCP09 tool [162]. Using a weather station outside of the centre is reasonable as the scope of buildings included in the data set includes many dwellings from the suburbs. Given the omission of flats from the sample set (which are more common properties in the city centre), this suburban weather data is likely to correspond more closely to the buildings simulated than central data.

The final file used in running EnergyPlus, the Input Data File (IDF), defines all of the attributes of the building itself, as well as simulation settings and heating schedules. Objects defined in the IDF are generally named, with the name acting as a unique identifier when referenced elsewhere. IDF manipulation was the primary method of specifying and updating buildings to be simulated in this research. The python module eppy was used to manipulate the IDFs pragmatically throughout this work [163]. Recently, a specialised open source tool has been released that provides an interface to perform energy simulations pragmatically [78]. This platform uses a similar method of IDF manipulation to the one used in this work, also using the eppy library to configure simulation settings. However, it was not available publicly when this research was underway.

A wall or roof is defined in the IDF as a building surface. Building Surfaces are defined by their construction (the layers or materials that make them up) as well as their position both in terms of physical coordinates and thermal zone membership. These relationships, and the relevant properties making them up, are shown in Figure 4.9. Surfaces are modelled like planes, defined by their vertices coordinates, with the thickness of the

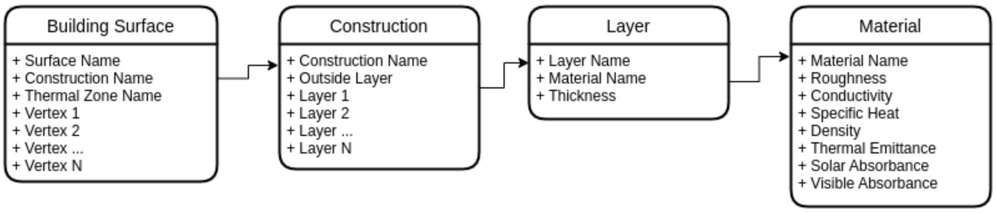

Figure 4.9: Attributes and relationships in EnergyPlus building surfaces.

third dimension derived from the properties of the construction making it up, which is further determined by the layer and material definitions in a hierarchical definition of properties. This is convenient for considering retrofits, as a retrofit can be performed on the model by substituting the construction name for an alternate construction without manipulating the physical dimensions of the surface.

Thermal zones are key to the granularity of the simulation process, as they are modelled as homogeneous with temperature loads averaged across them. When performing this work, thermal zones were assigned to each floor in a structure. This has been shown to provide a reasonable level of simulation accuracy [164] and will, for most households, approximate the living and sleeping zones used in the BREDEN method [60]. Physical attributes of a building, such as walls, floors, and ceilings, are assigned a thermal zone based on their geographic location. Thermal zones are also managed by their own heating system and can be assigned a schedule and setpoint.

Heating schedules represent the periods of time when a given heating system is active in a thermal zone. These are defined in the IDF and can be defined in terms of time, day, month and year. The heating setpoint defines the target temperature of a given zone and, while the schedule is active, the heating components in the zone will target it. The heating schedules and setpoints can be considered the only human behaviour modelled in the energy simulation stage of this research. While there are methods of integrating more sophisticated human behaviour modelling at this stage, such as window controlling behaviour, they are omitted from the scope of this work. This is partly a simplification, and partly because the objective function, heating energy demand, is likely to be minimally affected by window controlling behaviour.

The process of constructing IDF templates from the building data set, described in Section 4.5, required extrapolation. The two dimensional building footprints from the data set were projected into three dimensions using the eave and ridge heights. These structures were then divided internally using the number of floors, each of which was assigned a thermal zone. With a single thermal zone per floor, internal walls (the

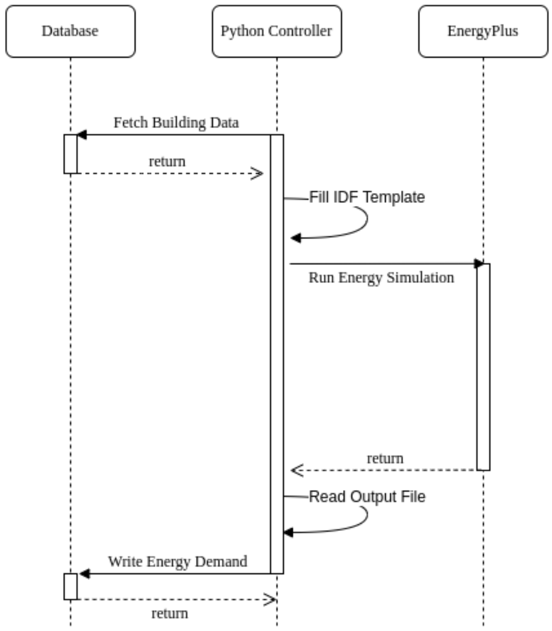

Figure 4.10: Sequence diagram showing a high level description of simulation of a single dwelling.

dimensions of which are unknown) were removed from the model scope. Windows were assigned to the building in accordance with the glazing ratio of a given side of the dwelling. Party walls of terrace houses are assumed adiabatic, based on the principle that temperature differences will be small and unknown when simulating individually. This also avoids any solar gains incorrectly attributed to the wall, precluding the need for adding shading objects to the simulation scene.

Control software, written in Python, was used to run the EnergyPlus simulation. The stages of energy evaluation for a non-retrofitted building are shown in 4.10. Later, when retrofits are performed, the base building IDF is altered in line with a given retrofit. The sampling and selection of individual retrofits, as well as the components within the scope to be retrofit, will be discussed when introducing the SEPM data set generation process in Chapter 5.

In order to efficiently construct IDFs for given buildings, a template containing the entire set of construction materials required in both the initial building stock and in the retrofit building versions was constructed. This allows the construction of surfaces using only the construction name required for a given value. The templates are filled using the controller, which performs the necessary extrapolations, geometric transformations, and simulation settings changes required to create a valid input file.

After the EnergyPlus simulation has been run, a set of output files are generated, reporting both on the results of the simulation and some metadata about the simulation process. Given the objective value used to train the SEPM will be annual heating energy demand, this value will be extracted and stored from simulation runs. The wall clock time, in seconds, taken to run the simulation will also be recorded in order to allow for speed comparisons and to allow some verification that simulations are run correctly. When simulations failed, usually due to an issue with the IDF, error reports are generated by EnergyPlus. Throughout the process of the research, all identified errors generated at this stage were corrected and the simulations re-run to ensure that no systematic bias in generated results was carried forward to the SEPM training.

Simulations were performed atomically, each considering a potential state of a single domestic dwelling, and thus excluding any shading effects, deemed reasonable given the low-rise nature of the data set used (average eave height 5.50m). As such, several EnergyPlus simulations can be run in parallel, with a single physical core assigned to each. Simulations were run using shell commands issued from the Python based supervisor program. Given the independent nature of the simulations, the output of the system grows linearly with the number of computer cores assigned, provided sufficient memory is available. System operations such as disk access were dominated by the simulation time for a given building, and so are not blocking.

4.6.2 Software Libraries

The primary development language used in this research was Python, as it benefits from a wide ecosystem of libraries to perform numerical and ML operations. There is also some evidence that python results in faster development time [165, 166]. Numerical operations and data pre-processing were performed using the Numpy and Pandas libraries. While Python is sometimes criticised for slow performance, Numpy uses an implementation in C with static typing and efficient memory management to perform operations up to 120 times faster than the equivalent speed in Python, rendering performance differences minimal compared with other languages [167]. This allows the developer to benefit from the features and ecosystem of Python while gaining the speed and performance benefits of lower level languages. Pandas is built on top of the Numpy data structures and provides additional functionality to allow convenient storage and processing of multidimensional and labelled data of heterogeneous types [168].

The Python ecosystem also includes a wide range of actively developed ML libraries. Keras and Tensorflow libraries were used for neural network formation, as well as some pre-processing tasks. The Keras Tuner library was used for hyperparameter tuning and model selection tasks. The Sci-kit learn library was used for other modelling tools, regression analysis and performance evaluation. The Matplotlib and Seaborn libraries were used for graphical visualisation.

As well as using Eppy [163] for Energyplus transformations, shapely [169] and geopandas [170] were used for geospatial transformations and data type support. Some visualisations were performed in qGIS [171] was used for some visualisation tasks, as well as for polygon simplification during pre-processing.

While Python and its ecosystem were used predominantly in this research, other statistics-focused libraries such as R and MATLAB would have been similarly suitable, with small advantages and disadvantages. The small speed and efficiency differences that may exist between software choices will be insignificant in the scale of speed differences achieved by the use of surrogate modelling itself, compared with the traditional approaches, and is therefore insignificant to the research.

Chapter 5

Modelling Purely Rational Agents Using Single-Objective Surrogate Optimisation

"He who would learn to fly one day must first learn to stand and walk and run and climb and dance; one cannot fly into flying."

Friedrich Nietzsche

This chapter will lay out the implementation and results of constructing a retrofit adoption model using a Surrogate Optimiser (SO) for rational, self-interested agents. The target system of the model will be the transformation, over time, of the residential building stock of the city of Nottingham. The construction of an initial Surrogate Energy Performance Model (SEPM) for energy demand evaluation is discussed and evaluated in Section 5.2. The process of optimisation using the SEPM combined with a Genetic Algorithm (GA) is laid out in 5.3. The optimisation data set is used to train the SO in Section 5.4, with the performance analysed against an unseen test set. Section 5.5 then applies the SO to the remaining test set and evaluates the quality and diversity of the predictions made before introducing a simple agent-based retrofit adoption model to see how the building stock may change over time. Section 5.6 will close with some points of discussion relating to the findings of the chapter and the utility of the method at this stage.

The initial model scope discussed in this chapter will assume rational and self-interested agents. The assumptions and methodological justifications of investigating agents of this type were discussed in Section 4.1. The assumption of self-interest in agents will be relaxed in Chapter 6 when extending the principle to multiple objectives.

5.1 Model Scope

When discussing model scope it is also worth considering the sub-components of the entire analysis, as there is a degree of modality among sub-components that make up the entire retrofit adoption model which is run in Section 5.5. However, when one sub-component is constructed from another, for example, when simulation models are used to train the SEPM, the assumptions, simplifications, and level of detail that make up the initial component are inherited. The interaction between these scopes is shown in Figure 5.1, which shows the outer scope comprised of the retrofit adoption model subsuming the attributes of other components used. This is demonstrated in Figure 5.2 which shows the stages of the process that must be revisited should a change in modelling scenario occur after the final analysis. The two key decision points occur during the simulation of energy demand (stage 2 in the diagram) at which point many of the assumptions of physical properties relating to the simulations are fixed. The second decision point occurs when performing the optimisation process, referred to as stage 4 in Figure 5.2.

Examples of assumptions or model inputs fixed after simulation stage:

- Retrofit options and their properties (e.g. insulation types, technology, thermostat types)
- Behavioural model for households (e.g. heating setpoints, heating schedules, window opening models)
- The output metric (e.g. heating demand, cooling demand, appliance models)

Examples of assumptions or model inputs fixed after Optimisation stage:

- Objective Values For agents (e.g. financial, environmental, multi-objective)

- Fuel price model
- Costs of different retrofit solutions (e.g. materials, labour, maintenance)
- Interest and discounting rates

These input parameters, modelling decisions, and assumptions become fixed at the point of evaluation (either in the simulation or optimisation stages). This holds for any encompassing scope that those output data are used in. This is to say that these exogenous variables become fixed assumptions for each given scenario after the simulation/optimisation stage in which they are used. While scenario analysis can be performed on these input variables/settings, or the later modelling stages be adjusted to create transparency of these inputs, they cannot be altered after the relevant stage without requiring recalculation and running of all later stages. These model parameters will be described for the simulation and optimisation procedures in Section 5.2 and Section 5.3 respectively.

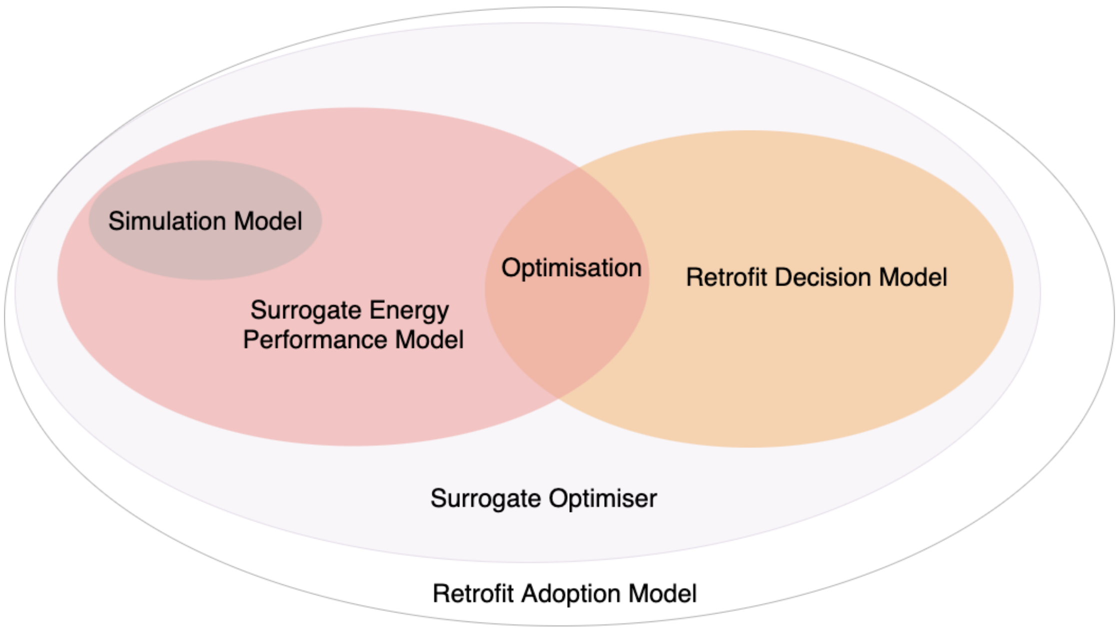

Figure 5.1: Venn Diagram showing the interaction between the scope and assumptions of sub-components.

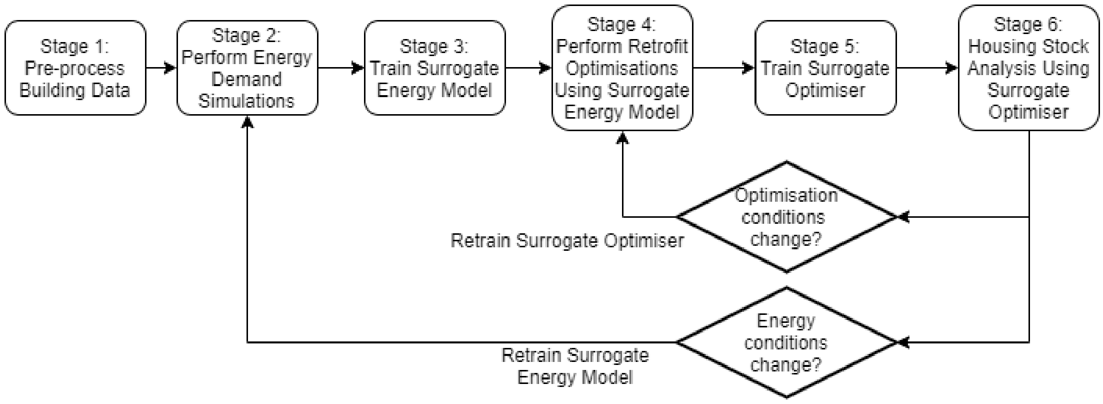

Figure 5.2: High level description of the method.

5.2 Surrogate Energy Performance Model Training

As discussed in 4.2, a SEPM will be used for energy evaluation in this research. In this section, the SEPM training process will be described, then the results presented. Given the objective value is the annual heating energy demand, expressed in gigajoules (GJ), a regression model is required for the analysis. While linear regression models were commonly used for this task in the past, recently Artificial Neural Networks (ANN) have become popular for the increased flexibility they offer, as discussed in Section 3.1.2. The process of training the ANN-based SEPM used in this study, including sampling, parameter tuning and evaluation are discussed in this section.

5.2.1 Retrofit choice selection

In order to analyse the retrofit decisions, a possible set of retrofit options must be considered. While there is often a focus on newer technologies in some research, a consensus among experts suggests that leading-edge technologies are less effective than simple physical retrofits such as wall and roof insulation [122]. When only considering heating energy demand, appliance modelling can also be unnecessary, as waste heat from appliances can be deducted from heating demand during scheduled heating periods [94]. As such, the retrofit choices in this initial exploration of the SO methodology will be kept limited to different physical methods, excluding smart appliances, automated shading, and other technological solutions that are still in relative infancy. This will include internal wall insulation, external wall insulation (such as cladding), window glazing thickness, and loft insulation. Various materials and thicknesses of these materials will be considered. The retrofit options chosen are presented in Table 5.1. The insulation thickness ranges were chosen to include the optimal thicknesses of the chosen materials under different conditions found across the retrofit literature [172, 173]. The thickness ranges were kept wide to account for the environmental benefits, which may exceed the economic ones [174]. This is likely to be considered more by the environmentally-conscious agents modelled in Chapter 6. Details of the physical embodied carbon properties of the insulation material included are presented in Table 5.2.

The heating system of dwellings will be considered to allow for both gas-based central heating systems and electric convection heaters. These are the two most common heat sources, representing 86% and 8% of national dwellings, respectively [175]. Retrofitting of either of these systems will be included within the scope of this research. While it is expected that NPV-optimising consumers would only consider retrofitting a gas central heating system, the inverse retrofit will also be considered. In a decarbonising grid scenario, it is possible that extremely emissions-sensitive households may consider an electrical retrofit from gas heating, as under the right scenario settings this may reduce net emissions.

The raw number of retrofit combinations possible for this optimisation problem is 18,000,000. This is a relatively small scale combinatorial optimisation problem when performed in isolation, however, each dwelling has unique characteristics which make them distinct problems. These asymmetries are reflected not just in the size and geometry of the dwellings but also in the existing fabric features (such as partially insulated buildings) which result in different objective function values. Households' carbon preferences will also play a role in altering objective function outcomes when environmentally conscious agents are considered. It is also worth noting that each solution evaluation requires an energy performance evaluation for a given building state, requiring the use of either energy simulation or a trained SEPM.

Component	Gene Name	Possible Values
External Wall Insulation Material	EWI_mat	Uninsulated, XPS, EPS, PIR
Internal Wall Insulation Material	IWI_mat	Uninsulated, XPS, EPS, PIR
External Wall Insulation Thickness	EWI_thick	30mm - 150mm in increments of 5mm
Internal Wall Insulation Thickness	IWI_thick	30mm - 150 mm in increments of 5mm
Heating Method	Heating	Electric, Gas Central Heating
Roof Insulation Material	Roof_Mat	None, Mineral Wool
Roof Insulation Thickness	Roof_thick	50mm - 400mm in increments of 25mm
Glazing Type	Glazing	Single Glazing, Double Glazing, Triple Glazing

Table 5.1: Description of possible retrofit components.

Material Name	Thermal Conductivity (W/mK)	Density (kg/mˆ3)	Emboddied Emissions (kgC02e/mˆ3)	Sources
Expanded Polystyrene (EPS)	0.029	29	3.29	[176], [177]
Extruded Polystyrene (XPS)	0.035	24	3.43	[176], [177]
Polyisocyanurate (PIR)	0.025	24	5.4	[176], [178]
Mineral Wool	0.04	20	1.12	[176], [177]

Table 5.2: Insulation material properties.

5.2.2 Data Set Generation Using EnergyPlus

The first stage of SEPM training requires the generation of a data set of energy performance results using the higher fidelity simulation technique, in this case, EnergyPlus. A detailed description of how EnergyPlus templates were generated and run can be found in Section 4.6.1, and a high level description of this data generation process is shown in Figure 5.3.

A random sampling method, with replacement, was used to obtain buildings from the data set to ensure a representative sample was selected. Replacement was used as the sample building can be retrofit many different ways and so multiple instances of the same building are not redundant in increasing sample diversity. An alternative method, Latin hyper-cube sampling, was also implemented. This form of stratified sampling subdivides the population into equally probable subsets before selecting a single random point from each subset [179]. However, the requirement for a known sample size and the need for an additional parameter was found prohibitive during model development. Given the relatively large sample size used and the relatively dense population data, naive random sampling was deemed sufficient. The final training set sample size was determined by the model performance gains which were beginning to diminish significantly at this size, as discussed in Section 5.2.4.

After the selection of a building from the data set, a random set of retrofit solutions were selected and applied. Each retrofit was selected independently for each building to ensure a diverse and representative set. Retrofits were applied by first constructing an EnergyPlus IDF of the building, then adjusting the appropriate physical components as discussed in Section 4.6.1. Simulations were then run and the energy demand was recorded. Verification was performed by monitoring the EnergyPlus output error files

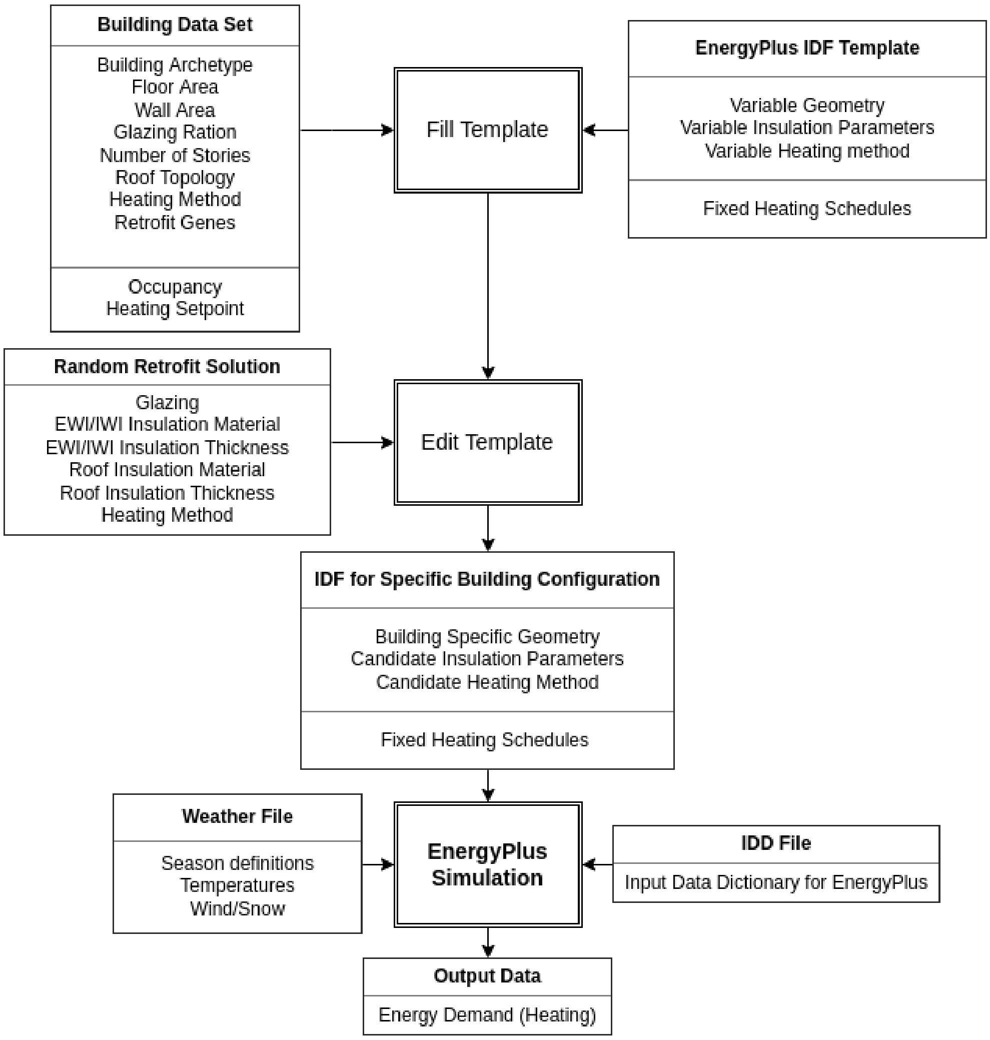

Figure 5.3: High level description of the EnergyPlus data generation stage.

(eplusout.err), with the detection of any errors with a severity rating above *warning* resulting in program termination and debugging.

5.2.3 EnergyPlus Simulation Results

The purpose of running the EnergyPlus simulations is to generate a data set for SEPM training. The validity of these simulations, which is to say the closeness of the outcomes of the models to the target systems they represent, is of relatively minor importance for the methodological study provided the complexity of the simulation system it presents.

This is especially true as both the physical simulation process itself and the generation of SEPMs from its results are well established as valid in the literature (See Section 3.1.1. As such, a simple statistical validation technique was used to ensure the distribution of simulation results broadly matched the distribution of heating demand values found for alternative sources.

The median national space heating energy consumption in the UK was reported at 10,118 kWh per household in 2019, an equivalent of 107 kWh per m^2 [180]. These average values are slightly higher than the simulated, non-retrofitted simulation values of the urban building stock sample of 10,030.8 kWh per household and 103.8 kWh per m^2. Given the national figures were generated from a nationwide sampling methodology rather than a bottom-up simulation methodology, this variation in median outcomes of approximately 1% and 3% in per dwelling and per unit area energy usage was deemed to be within an acceptable margin. The distribution of heating energy demand values is in line with the most recent whole stock energy distribution data found, although the data is from an older 2009 survey in which the median energy consumption was higher [181]. The distribution of existing stock annual heating energy demand is shown in Figure 5.4.

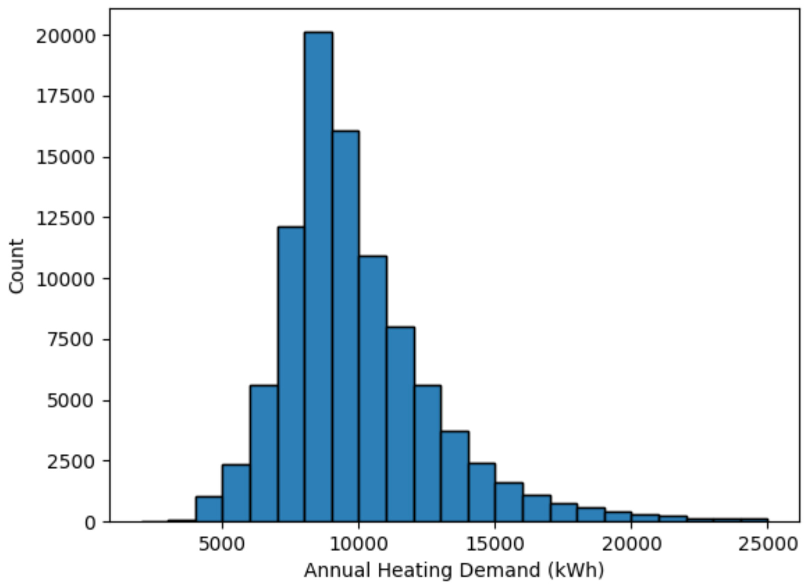

Figure 5.4: Distribution of simulated household space heating energy demand (kWh/year).

5.2.4 Surrogate Energy Performance Model Training

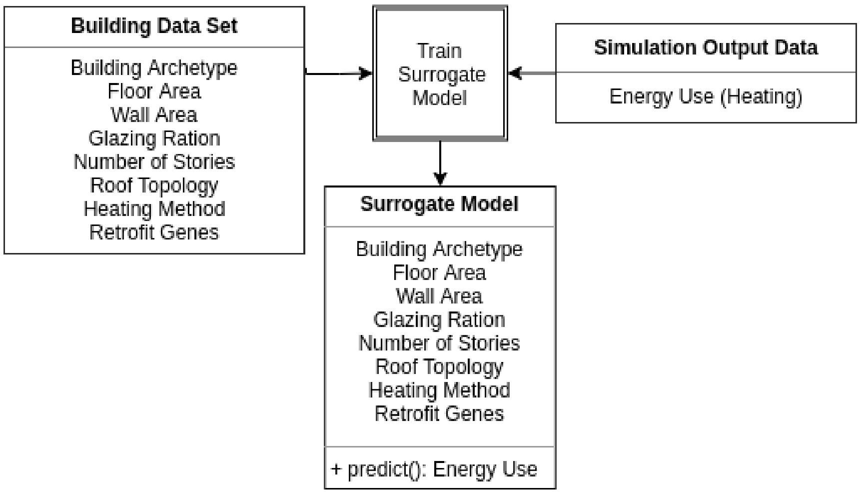

Figure 5.5: SEPM training process using the energy performance data set.

5.2.5 Model Hyperparameters

The model training parameters were tuned using a grid search method, with 20 repetitions per settings combination, scored using the mean r^2. The grid settings can be found in Table 5.4. The selection of model hyperparameters for the grid search was taken and adapted from ANN-based SEPMs found in related works, which are discussed in Section 3.1.2. Some of these were adapted to match the shape of the training data used. It should be noted that while the ANN was tuned, the performance of the model was only minimally affected by the hyperparameter section process, and although most parameters produced a statistically significant difference in model performance with the samples used, the absolute difference in performance is minimal due to the high performance of this model. This is likely due to the relatively large sample size combined with the inclusion of grid settings which had been found to perform well in related works. An Ordinary Least Squares (OLS) analysis on the Root Mean Square Error (RMSE) of the sample of trained models is shown in Table 5.3.

Dep. Variable:	rmse	**R-squared:**	0.133
Model:	OLS	**Adj. R-squared:**	0.123
Method:	Least Squares	**F-statistic:**	12.68
Date:	Fri, 26 Nov 2021	**Prob (F-statistic):**	1.79e-11
Time:	07:06:30	**Log-Likelihood:**	-50.149
No. Observations:	418	**AIC:**	112.3
Df Residuals:	412	**BIC:**	136.5
Df Model:	5		

	coef	**std err**	**t**	**P> |t|**	**[0.025**	**0.975]**
const	1.3714	0.028	49.271	0.000	1.317	1.426
network_[128, 64, 32]	-0.1194	0.041	-2.902	0.004	-0.200	-0.039
network_[16]	0.1648	0.041	4.037	0.000	0.085	0.245
batch_size_8	-0.1268	0.038	-3.343	0.001	-0.201	-0.052
batch_size_16	-0.1876	0.038	-4.945	0.000	-0.262	-0.113
batch_size_32	-0.1958	0.038	-5.163	0.000	-0.270	-0.121

Omnibus:	269.175	**Durbin-Watson:**	2.171
Prob(Omnibus):	0.000	**Jarque-Bera (JB):**	2945.978
Skew:	2.629	**Prob(JB):**	0.00
Kurtosis:	14.895	**Cond. No.**	4.85

Notes:
[1] Standard Errors assume that the covariance matrix of the errors is correctly specified.

Table 5.3: Results of OLS analysis on parameter sweep of SEPM model after iterative removing statistically insignificant effects (p¡0.05).

The loss function, used to evaluate the model during training, was least squared errors (L2) as this is preferred for models not containing a large number of outliers. The Rectified Linear Unit (ReLU) activation function was used on all layers to provide the desired non-linearity in the neural network training. The Adam optimisation function was selected as a fast, out of the box training function, precluding the requirement for manual learning rate selection [182]. These fixed settings were selected as sensible defaults to reduce the parameter tuning space and correspond to standard choices in the literature for ANN-based SEPMs [73, 74, 82]. To determine the number of desired training epochs, the validation loss was measured against the training loss to ensure no major divergence. As can be seen in Figure 5.6, there is little evidence of over-fitting with training and validation loss converging and the variance of validation loss diminishing approximately 500 epochs. As such, the maximum training epochs was set to 1,000, with an early stopping procedure implemented. Training was automatically stopped if validation loss stalled below 0.0001 over a patience of 50 epochs.

The selection of network topologies for grid search was based on adapting the successful topologies used in related works (see Section 3.1.2) and adapted to input parameter size [183, 81, 96, 184]. Both the number of layers and the number of nodes in each layer were altered to test for changes in model performance. The best performing network topology comprised three hidden layers of decreasing size (128, 64, 32). The increase in performance compared with other topologies was small (-0.12 RMSE) but statistically significant at the 1% level. In comparison, the worst performing topology was a single layer of 16 nodes, indicating that the additional network nodes and layers were useful to model training.

The pre-processing methods selected for grid search were normalisation, standardisation and logarithmic transformation. The purpose of pre-processing continuous input variables is to remove scale effects. Having some inputs, such as total area, with values of several hundred, which other variables such as glazing ration (expressed as decimals) can lead to inefficient training and a loss of accuracy. This scaling effect can be mitigated by pre-processing methods. Normalisation forces all values into a scale between zero and one using min-max scaling, standardisation converts values in a sample to be expressed in terms of the number of standard deviations that value lies from the mean of the sample, and logarithmic transformation places input variables on a logarithmic scale, reducing the impact of scale without completely removing it. The choice between these pre-processing methods was statistically insignificant during the parameter tuning grid search.

Setting	Grid Values
Model Type	Regression
Pre-proccessing	Normalisation, Logarithmic, Standardisation
Hidden Layers	1/ 16, 32, 64
	2/ [64, 32], [32, 64]
	3/ [64,32, 16], [128, 64, 32], [32, 64, 128]
Activation Function	ReLU
Loss Function	L2 Loss
Batch Size	16, 32, 64
Training Function	Adam

Table 5.4: Description of SEPM hyperparameters. Grid search parameters are separated by commas.

5.2.6 Sample Size

Given the SEPM training set is generated using a simulation procedure, the upper limit for sample generation is bound by the computational time for simulation and model training, rather than real-world data constraints. As such, the effect of training set size on model performance was tested to determine an appropriate trade-off between model performance and computational time. The sample size was incremented in small intervals of 500 between 0 and 10,000, after which it was incremented in intervals of 5,000 up to a total of 140,000. Model training was repeated 20 times per sample size, for a total of 920 evaluated models. The effect of sample size on the model's r^2 value are shown in Figure 5.7. An assessment of the performance data revealed a diminishing return on sample size, with even small training sets of 5,000 producing a mean r^2 value of 0.988, increasing to 0.995 for the largest sample sizes. The diminishing nature of the performance returned was demonstrated visually by fitting a trend line of the form $y\ log(x)$ through the points, with the sample sizes stratified to those below 5,000 and those above to distinguish the effect between small and very large samples. A Pearson correlation coefficient of 0.416 (p<0.01) confirmed a positive correlation between sample size and model performance across the range. This held when smaller sample sizes (n<20,000) were excluded from the sample, with a correlation coefficient of 0.361 (p<0.01), showing that the performance continued to increase, albeit at a decreased rate, across the range. Given the positive but diminishing returns to sample sizes, a large sample of 140,000 was used for models in later stages of the research, while smaller samples sizes were used in parameter tuning experiments in which model training time represented significant computational cost.

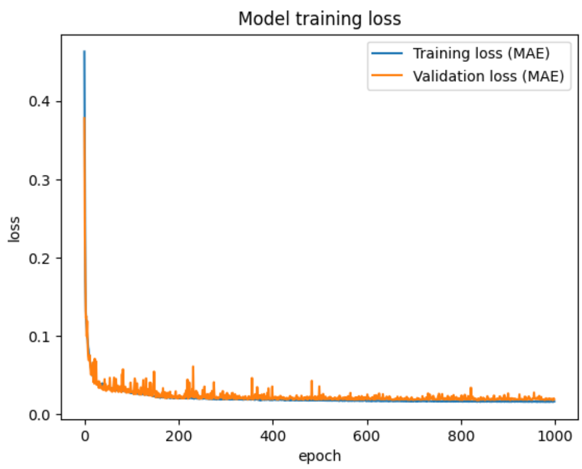

Figure 5.6: The effect of the number of training epochs on training and validation loss.

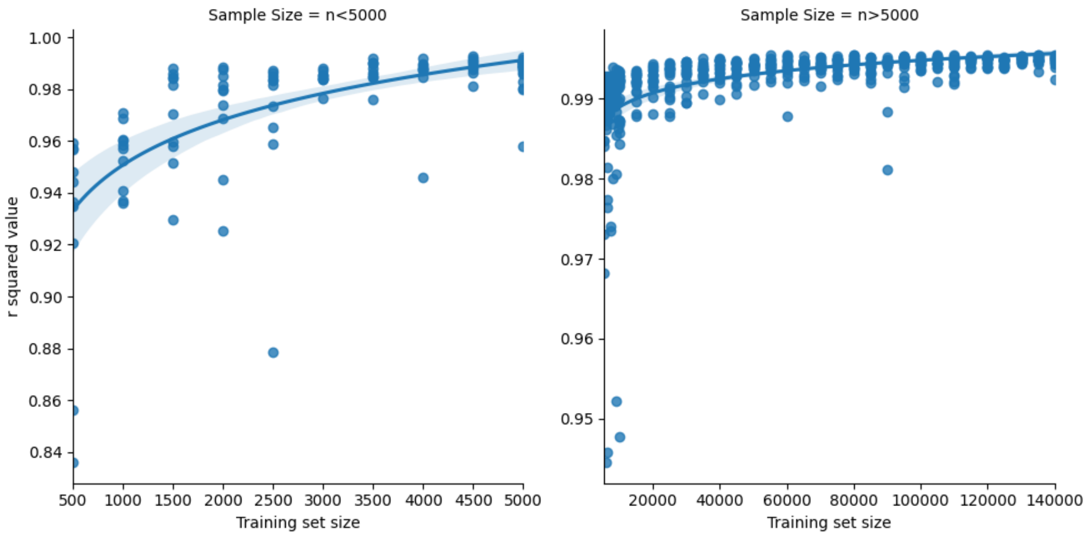

Figure 5.7: The effect of training set size on SEPM performance. Trend line fitted using $y\ log(x)$ linear regression.

5.2.7 Feature Selection

Alongside the candidate features from the initial data set, some aggregate geometric features are derived from the EnergyPlus models. These included the front, back and side glazing ratios of the building, the floor area, the number of stories, and the orientation of the property. Household details such as the temperature setpoint and heating schedule were also included as key behavioural variables, which are known to have a significant impact on energy demand [183]. Details about common SEPM features and their relative significance can be found in the related works of Section 3.1.2

A backwards feature selection process was performed, with features eliminated from the model systematically and the model retrained. A sample of 50 models was trained for each removed variable, and a Student's t-test was performed on the RMSE to determine if the model performed worse with a given variable removed. The results of the final composite features are shown in Table 5.5. Given they are required for the optimisation process and they represent a large number of important features, the inputs corresponding to model retrofits themselves were excluded from the backwards feature selection process. The selected features are generally related to the physical properties of the building, with the most significant physical feature being the total area of the building, followed by the wall area and wall type. The least practically significant physical feature in terms of additional error was the glazing area of the building. This could be due to the relatively low glazing ratio of the urban domestic stock compared with retail or commercial buildings where glazed façades are more common.

Removed Attribute	Mean RMSE (GJ)	Mean r2	Mean Absolute Error (GJ)
None	1.162693	0.990426	0.690875
GlazeArea	1.204369	0.989328	0.719256
rTYPOLOGY	1.213802	0.989360	0.728794
Storeys	1.217752	0.989376	0.725345
BHA_relh2	1.230255	0.988999	0.704932
crAGE	1.235229	0.988668	0.740531
crTYPE	1.287131	0.988020	0.764233
WType	1.355030	0.987092	0.797539
WallArea	1.446104	0.984717	0.864996
TotalArea	2.920386	0.939051	1.807114

Table 5.5: Results of backwards feature selection during SEPM training shows remaining features have statistically significant effect on model RMSE when removed.

5.2.8 SEPM Results

The final model, which will be used going forward at the optimisation stage, was initialised with random weights and was trained using a batch size of 32 for 247 epochs before being stopped automatically due to a stalling validation loss. The model was trained using 70% of the data, with the remaining 30% split evenly between validation and test sets from an entire sample size of 140,000. The RMSE was 0.864 GJ while the adjusted r^2 was 0.995. The mean absolute percentage error of the trained model when applied to the test data was 1.59%. These metrics indicate a well performing model in aggregate. A scatter plot showing the simulated annual energy use against the SEPM prediction is shown in Figure 5.8, demonstrating relatively few outliers. The slight fanning out of values indicates an increase in the magnitude and variance in error for higher demand buildings, although this effect is not large.

An analysis of the SEPM residuals indicated that they were normally distributed. This can be seen in Figure 5.9 and was confirmed using a Shapiro-Wilk test ($p < 0.01$). The distribution of errors is important for the optimisation process, as a biased SEPM may lead to, for example, overestimation of non-retrofitted building energy demand which would lead to overestimation of retrofit benefits. As such, the errors of the individual retrofit solutions were analysed and can be seen for the insulation materials in Figure 5.10 and for the different thicknesses in 5.11. There is no indication that the magnitude or distribution of errors varied across the range of possible retrofit solutions, allowing the model to be deployed in performance analysis of the retrofit measures modelled.

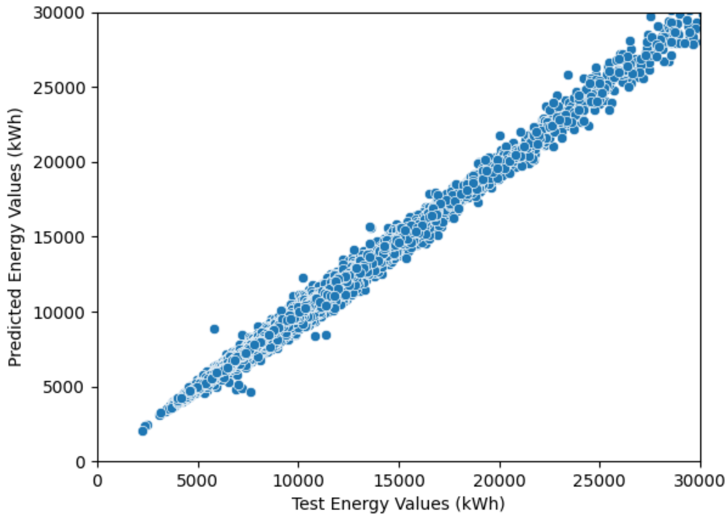

Figure 5.8: Scatter plot of simulated annual heating demand against the SEPM predictions of the test set.

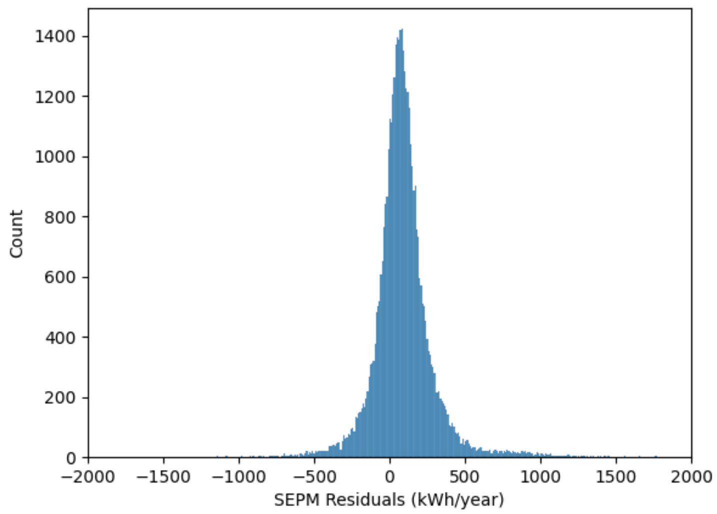

Figure 5.9: Histogram of SEPM test set residuals, showing a normal distribution.

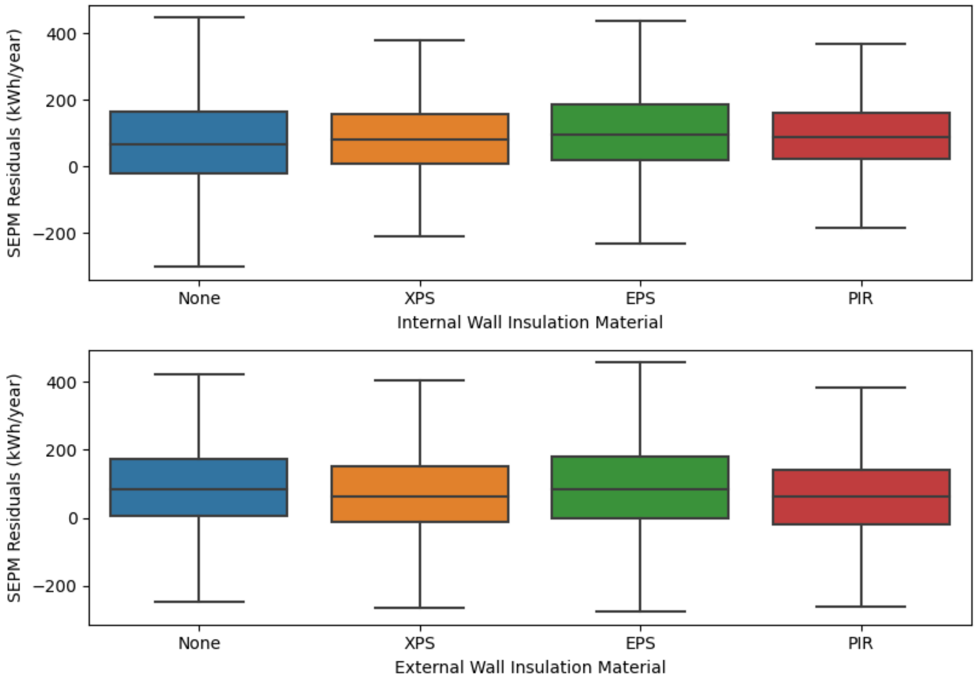

Figure 5.10: Box plot of the retrofit insulation materials impact on the SEPM residuals, showing no significant heteroskedasticity.

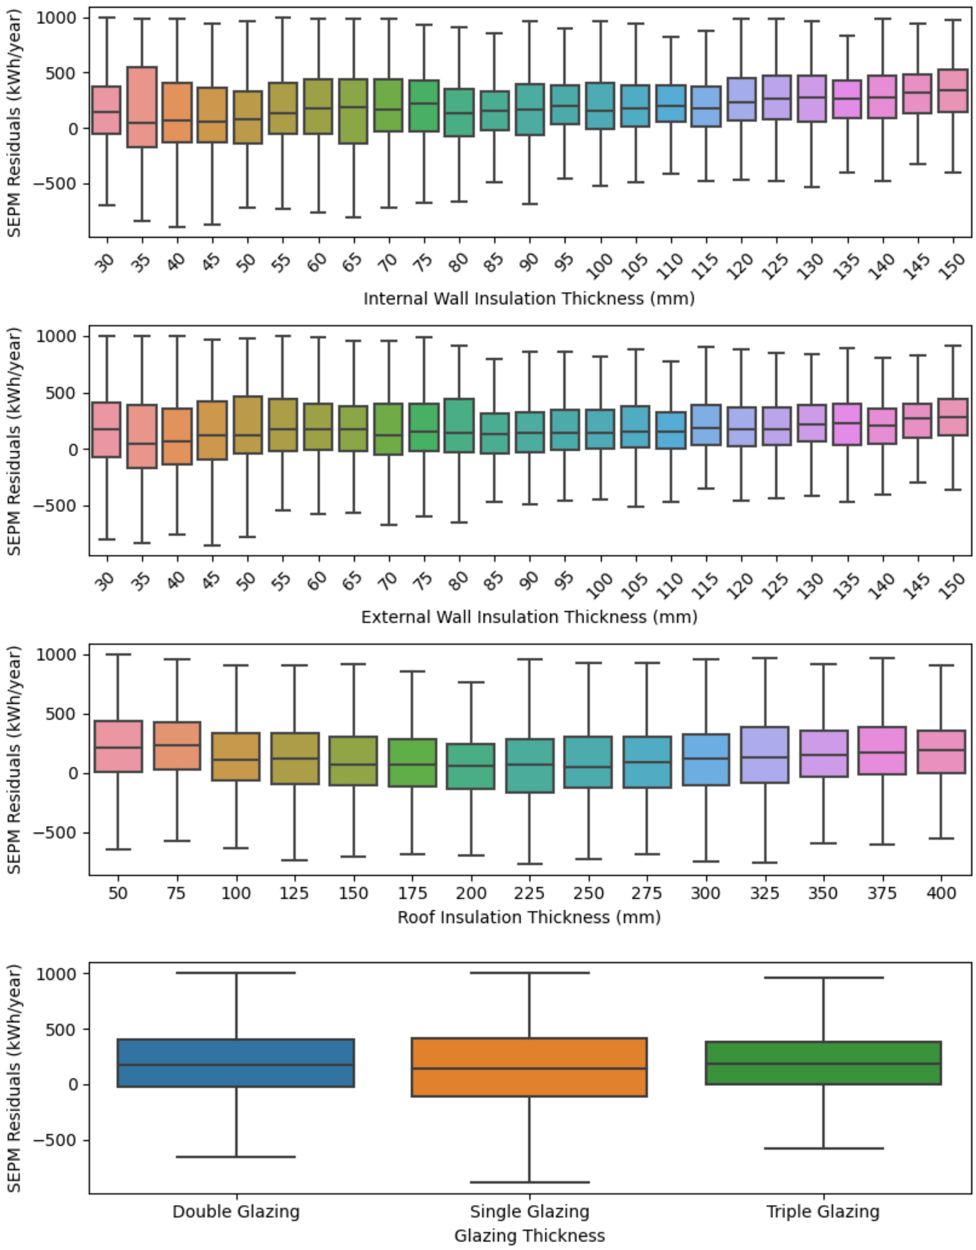

Figure 5.11: Box plot of the retrofit insulation thickness impact on the SEPM residuals, showing no significant heteroskedasticity.

5.3 Single-Objective Optimisation using a Genetic Algorithm

The methodology of optimisation using SEPMs is discussed in Section 4.3, with related work found in Section 3.1. A GA was used to obtain the required set of near-optimal solutions.

5.3.1 Candidate Evaluation

5.3.1.a Agent Design

The objective function used for the self-interested agents, the NPV of the investments, can be found in Equation 5.1. The decision model for agents of this type, given the assumptions laid out in Section 4.1, is quite simple. First, WHRSs for the dwelling are optimised with respect to NPV. If the NPV-maximising solution has a positive NPV, then the agent is better off installing the retrofit than not. This decision model is slightly different to the full retrofit adoption model, as these agents have already decided to evaluate possible retrofit solutions. The full retrofit adoption model will consider the rate of solution evaluation over time, as well as account for triggers and constraints that cause or prevent the installation of a retrofit by different households.

$$NPV = C_0 - \sum_{t=0}^{n} \frac{R_t}{(1+i)^t} \tag{5.1}$$

Equation 5.1: NPV for an investment of C_0, returns of R at time t and a discount rate of i.

As a metric of evaluation used by the rational and self-interested agents considered, NPV is quite simple in principle. However, the process of evaluating candidate solutions requires several stages of calculation. A description of the evaluation process, including the parameters of the relevant models, is shown in Figure 5.12. The SEPM is initially called to ascertain the baseline energy demands for a given dwelling. The model is then called once per candidate evaluation, using the candidate solution building properties, to determine the change in energy demand compared with the initial building state. The savings and cost models, which are outlined below, then assess the initial and ongoing financial information required to calculate the NPV used to score solutions.

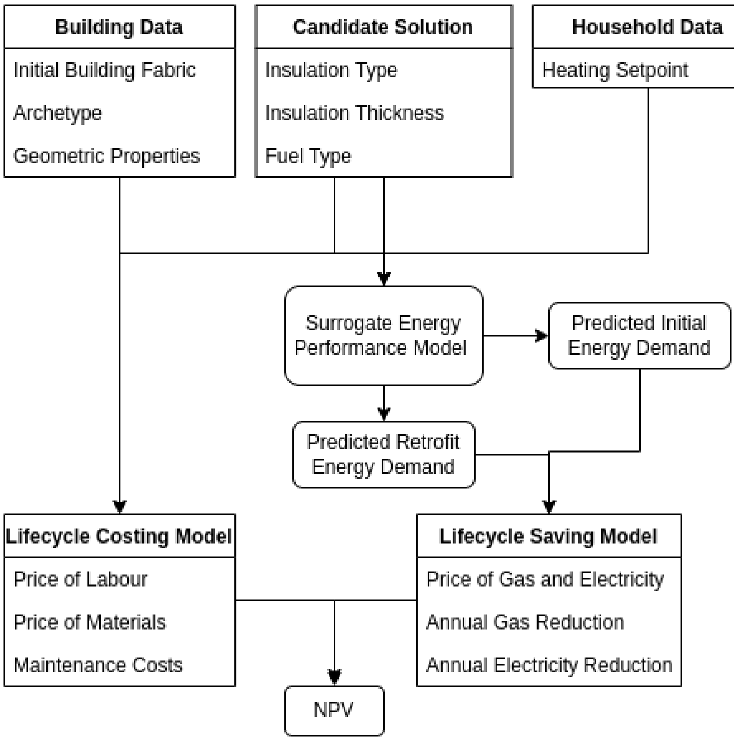

Figure 5.12: Description of a candidate solution's NPV calculation including exogenous variables.

5.3.1.b Financial Cost Model

The installation cost of retrofits was calculated based on the intervention chosen. The cost model implemented decomposed retrofit costs into material and labour. Insulation was costed on a cost per m^2 rate, while glazing and gas boilers were costed on a per unit basis. Material costs were sourced using a range of commercial material providers with some interpolation where a material thickness was not available from a given supplier [185, 186, 187, 188]. Full retrofit costing tables are laid out in Section A.2.1 of the Appendix.

Costs are calculated for each component individually and summed to obtain the total cost, without accounting for possible economies of installation scale (e.g. from simultaneous retrofit of a row of terraced housing). This is justified given the household scope of the adoption model, excluding housing associations and city councils that are able to access such economies. When interviewed by researchers for the department of business, energy and industrial strategy, installers reported no significant scale economies in domestic retrofits for glazing, wall or boiler retrofits [1].

Installation costs are also modelled as averages per m^2 for a given material type of thickness or per unit for glazing/boiler components. These averages attempt to capture, in the labour component, additional confounding cost factors such as detailing requirements or reparation of thermal bridges created. The costing model is designed to be modular, allowing for more granular costing models should more granular building data become available, or if they are required for a certain use case. The inputs can also be considered parameters instead of fixed values, allowing for analysis based on forecast cost scenarios.

Retrofit Type	Material Name	Labour Cost	Model Mat. Cost Min	Max	Validation Cost Min	Mid	Max	Validation Loss Min	Mid	Max
EWI	XPS	100	114.50	127.88	116	148	180	1.50	-26.81	52.12
EWI	EPS	100	111.67	120	116	148	180	4.33	-32.17	60
EWI	PIR	100	112.99	120.58	116	148	180	3.01	-31.22	59.42
IWI	XPS	85	96.67	105	95	117.5	140	-1.67	-16.67	35
IWI	EPS	85	97.99	105.58	95	117.5	140	-2.99	-15.72	34.42
IWI	PIR	85	97.99	105.58	95	117.5	140	-2.99	-15.72	34.42
Roof	Mineral Wool	30	33.26	46.09	20	30	40	-13.26	9.67	-6.09
Heating	Gas Heating	3,000	3,000	3,000	2700	3250	3,800	-300	-250	800
Glazing	Single	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Glazing	Double	50	375	375	300	475	650	-75	-100	275
Glazing	Triple	50	475	475	300	475	650	-175	0	175

Table 5.6: Validation source [1]. Values are in GBP (£) per m^2 for insulation, per window for glazing, and per dwelling for heating.

The financial costs model was validated against the costs assumptions laid out by the retrofit cost report from the Department of Business, Energy and Industrial Strategy [1]. It was decided to use a labour-and-materials based cost model, instead of relying on externally aggregated figures, to give the modeller additional control over model parameters. This allows for scenario analysis such as a labour shortage or material subsidies for the purposes of a retrofit that would not be possible using scaled aggregates. The results of the validation are shown in 5.6, demonstrating that the model provided reasonable estimates for the lower to mid-tier cost estimates but performed poorly in capturing the high-end of the scale. This is likely due to the survey-based methodology used in the aggregate cost reports, which is likely to be skewed at the extremes due to anomalous features in a minority of dwellings. This is shown by the significant cost variance present in the survey reports and acknowledged in the methodological notes.

5.3.1.c Financial Savings Model

The financial savings of the retrofits considered are derived from either energy savings, which result in reduced fuel consumption, or a change in fuel source from a more expensive to a cheaper supply. The NPV of savings (S) in a dual fuel dwelling over the lifecycle of a retrofit can be calculated as shown in Equation 5.2, with f_{et} representing the reduction in the demand for electricity at time t (with $t = 0$ representing the current year), p_{er} the projected price per unit of electricity at time t. f_g and p_g similarly represent the fall in gas demand as a result of the retrofit, and the projected price of gas at time t.

$$S = \sum_{t=0}^{N} \frac{f_{et}p_{et} + f_{gt}p_{gt}}{(1+r)^t} \tag{5.2}$$

Equation 5.2: NPV of Energy Savings over N years with a discount rate of r. The price of electricity and gas in year t given as p_{et} and p_{gt}. The change in demand (for year t) of electricity and gas as a result of the retrofit as f_{et} and f_{gt}.

Variable i represents the discount rate. This rate was set to approximate the risk-free rate of return in the UK, with 0.5% used in baseline scenarios. While historically this has been above 2% [189], recent yields have plateaued to below 1% [190]. The exact discounting rate used would depend on the interests of the modeller and the framing desired. Risk-free rates of return were used to model relatively rational agents acting in a liquid economy. If the modeller is dealing with highly risk-averse populations or in an economy with low liquidity, such as after the 2008 financial crisis, a higher discounting rate may be considered. It would also be possible to apply heterogeneous discounting rates for different agents based on their preferences. However, this is considered outside of the scope of this research as the data is unavailable and the option does not relate to the research question at hand.

Fuel prices were sourced from the BRE group, the division responsible for the SAP methodology used across government departments [191][1]. However, given that only current fuel prices are provided, a method of determining future fuel prices is both necessary and desirable for the analysis of the financial performance of a given energy-saving retrofit. In this implementation, an inflation rate of 3% was applied annually

[1] The rdSAP energy prices from January 2020 were used for the majority of this research. However, rdSAP values have not significantly changed above inflation over the period of study.

for fuel prices to approximately follow the recent SAP price trends. However, this approach is exogenous and more sophisticated fuel pricing models could easily be applied if necessary. The modular fuel price model could also allow for extreme scenarios such as a carbon tax which, under a decarbonising electricity grid, may result in the costs of electricity falling below gas in some cases.

5.3.2 Genetic Algorithm Design

A Genetic Algorithm (GA) was used for the optimisation procedure in this work. As discussed in Sections 3.1 and 3.2.3, GAs are a popular and robust choice for problems in these domains and have commonly been combined with SEPMs for evaluation. The optimisation stage can be performed using any meta-heuristic, provided the near-optimal retrofit solutions represent decisions made by the households modelled in the optimisation decision stage. A background discussion of GAs and their components can be found in Section 2.3 alongside a description of alternative optimisation meta-heuristics.

In order to tune the GA settings, a parameter grid experiment was performed. The list of settings can be seen in Table 5.7. The results of this parameter experiment will be discussed in this section and the aggregate results can be found in Section A.2.2 of the Appendix. The tuning was performed on a sample of 300 buildings selected using random initial building sampling with a grid size of 144 variable combinations. The optimisation was replicated 30 times for each building-setting combination for a total of 1,296,000 optimisation runs. The mean NPV of the replications was used as a performance metric, with the variance of NPV values also recorded to determine the consistency of a given setting's performance. A description of the settings tested, alongside the final selection based on the settings' experimentation, can be found below.

Setting	Values
Stopping Condition	Max Calls
Max Calls	500
Mutation Method	Uniform, Random
Mate Selection Method	Tournament
Tournament Size	2
Recombination Method	2 Point, Uniform
Replacement Method	Pure Elitist, Soft Elitist
Mutation Rate	5%, 10%
Objective Function	NPV
Population Size	8, 16, 32

Table 5.7: Settings grid used for Genetic Algorithm parameter tuning grid.

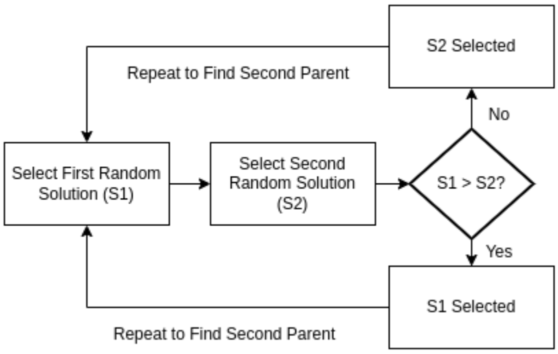

Figure 5.13: Flow chart of tournament selection process with a tournament size 2.

Mates were selected using a tournament mate selection method, as this was found in the literature to be a robust choice for optimisation when using a surrogate model [192, 193, 194]. The tournament selection process involves the selection of n solutions from the solution pool at random. These solutions then compete, with the best performing selected for reproduction. This is repeated to obtain the second parent. A flow chart of this process can be seen in Figure 5.13 using a tournament size of two. Tournaments of size two were used with replacement, meaning a solution could reproduce multiple times per generation.

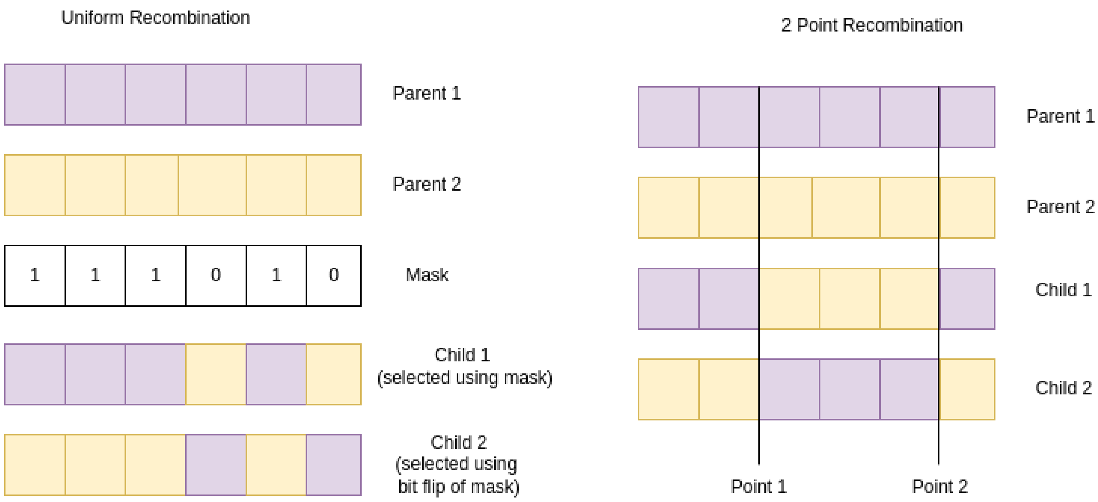

Figure 5.14: Demonstration of 2 point and Uniform combination methods.

The recombination method, sometimes referred to as crossover, determines how two selected parents are recombined to create offspring. A uniform recombination strategy applies a binary mask to the length of the genome, selecting each gene with equal probability from one parent of the other. This creates two children by applying the bit flipped inverse mask for the second offspring. In contrast, a 2 point recombination method selects two locations within the parent genes, cutting them and splicing them at those points into two offspring. A diagram demonstrating the effect of the different combination methods is shown in Figure 5.14. The result of the parameter experiment showed that the selection of the recombination method was not significant, a finding in common with related literature [193]. However, average performance was slightly higher using 2 point recombination, and the variance of the 2 point recombination method was also lower, suggesting more consistent performance.

The mutation method is responsible for altering child genes, increasing the diversity of solutions in the set and giving the chance to escape local maxima. Mutations occur at a given probability, the mutation rate, which is applied independently to each gene in the population. The random mutation method implemented replaces prior values of mutating genes with a random integer in the valid range. In contrast, the uniform mutation method performs a local change, increasing or decreasing the gene value by one and, therefore, causes less diversity in the gene pool. The more exploratory random-mutation method performed best, resulting in a higher mean NPV and lower NPV variance. Mutations were also directed away from dominated spaces to reduce computational speed; for example, replacing existing insulation with thinner substitutes was not considered viable.

The stopping condition was set to a maximum number of function calls, which was set to 500 retrofit evaluations. An additional stopping condition was in place if the quality of the best performing solution does not improve by a given relative tolerance level over the course of a number of stall generations. When optimising real solutions the tolerance level was set to 5e-5 and the stall generations were fixed to 6 to detect when the solutions had reached a local maximum.

The replacement method determines which solutions from the parent and child generations are replaced before the next mate selection phase. An elitist replacement strategy ensures that focus is placed on the best performing solutions, adding selective pressure. Two alternative elitist candidate strategies were tested. The pure elitist setting involved the selection of the best performing solutions from both the parent and child generation, with the worst performing half replaced. This method places significant selective pressure. In the soft elitist setting, only the single best performing parent is selected, with the rest of the parents replaced by the child population. This strategy maintains a greater diversity of solutions while potentially sacrificing convergence speed. The best performing strategy was the pure elitism method, although the difference was not large, as shown in Section A.2.2 of the Appendix.

The population size represents the number of solutions retained after each replacement procedure occurs, representing the size of each generation. Original population sizes were determined proportionally to the size of the genome length as 8, 16 and 32 as small population sizes have been found to perform well on similar problems [193]. The initial genome found that population size the most significant setting from the grid search, with populations of 32 and more resulting in better mean performance and smaller performance variance. Given the beneficial impact of increased population size, a second parameter tuning experiment was performed with increasing population sizes. The other settings were fixed to their determined values to reduce the number of repetitions required. The performance increased until a population size of 48, at which point the mean NPV of solutions plateaued. As such, population sizes of 48 were used, as this minimised the wall clock optimisation time over the larger population size while maintaining the same solution quality.

5.3.3 Genetic Algorithm Optimisation Results

The results of the traditional optimisation process, as well as being interesting as a sample of the population of the stock, also represent the input data for the SO used in Section 5.4. As such, the data will be investigated and visualised for the detection of patterns. Solutions will be referred to as WHRSs, which are solutions that have a positive utility (in this case, NPV) when implemented on a given building. Solutions are specific to the properties of the building in question and are comprised of one or more of the retrofit installations options discussed in Section 5.2.1. A sample of 10,000 buildings was used to gain insight into the distribution of the stock and for training in later stages.

An important note regarding the results of the optimisation stage is that the results are specific to the optimisation scenario parameters as well as any assumptions and simplifications made in the model design. While the physical properties of the retrofits are dependent on the SEPM results, and so are essentially fixed at this stage, the impact of those physical properties changes on the objective function and, therefore, optimal installations are input parameters that have been laid out. Given the NPV objective function, the most significant parameters are likely to be the costs of the retrofit installations, the prices of fuel sources, and the discounting rate, as these make up the primary components of the NPV calculation for a given energy saving.

The proportion of the stock sample for which an NPV-positive WHRS was 37.4%. The average WHRS investment was £7,112, with an average NPV of £5,148, indicating total

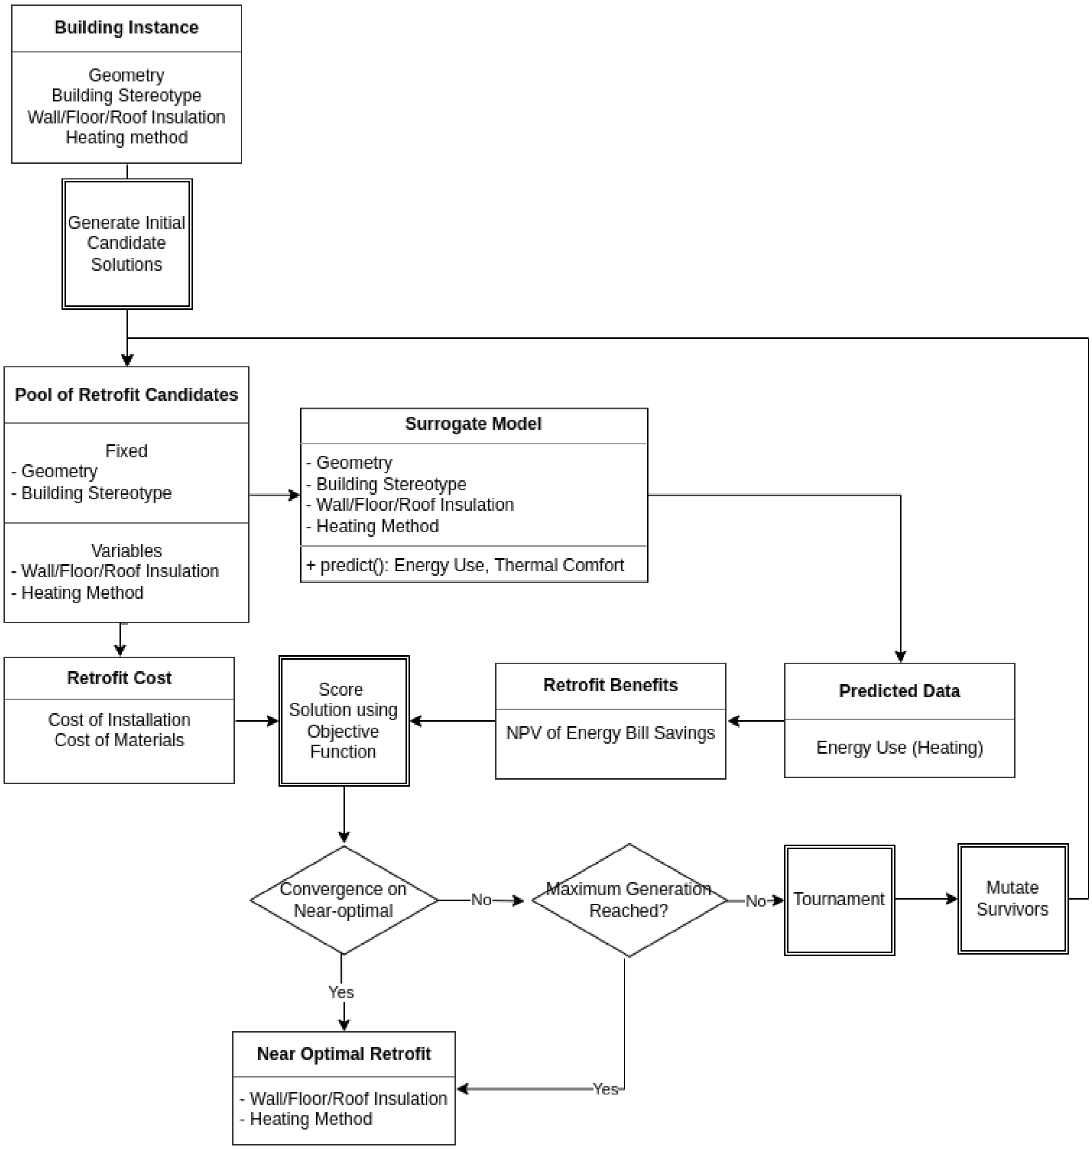

Figure 5.15: Process flow of the retrofit optimisation process.

present-value energy bill savings of £12,260 over the lifecycle of the installation. The average annualised ROI was 2.86%, although this figure already includes any opportunity cost conceptualised into the discount rate, resulting in higher real returns. The majority (79.28%) of WHRSs were single measures, indicating that a fair few of the dwellings' retrofits could profitably be installed with multiple measures for the scenario parameters used. However, 18.8% of the WHRSs did include two measures and in 1.7% it was found that three were optimal, showing the value of finding WHRS combinations rather than examining each potential measure in isolation. The average net lifecycle emissions reduction was 28.31 tCO_2e[2].

Wall insulation was the most popular optimal installation, with 27.5% of all buildings finding a WHRS which included wall insulation, 73.22% of all the retrofits found. This is likely due to the low rates of additional wall insulation in the base stock, with only 2.96% of the existing stock having any additional wall insulation[3]. In comparison, 90.5% of the existing stock had some form of loft/roof insulation already, leaving less potential for dramatic retrofits. Given the relatively high cost of wall insulation, the average WHRS investment was highest when this measure was included, averaging £8,771 per WHRS. The average return was also lower, with the NPV averaging £3,099, and an effective annualised ROI of 1.41%. This low ROI is likely to make this retrofit more sensitive to discounting, as highly discounted future energy bills are less likely to make a positive net return on a high initial investment. However, the significant financial investments in these retrofits do result in large energy reductions, averaging 4,988.1 kWh/year and a net lifecycle emissions reduction of 36.07 tCO_2e. The majority of optimal wall insulation retrofits under these optimisation scenario settings were internal installations of Polyisocyanurate (PIR), with the remainder being expanded Expanded Polystyrene (EPS). The dominance of internal installations when NPV optimising is unsurprising in this case, as the costs are higher when placing equivalent physical materials due to additional labour costs associated with external wall insulation.

The breakdown of wall insulation installations showed that two story terrace houses were over represented in the sample of WHRSs compared with the population (19.2% vs 6.3%). This is to be expected, as the adiabatic walls modelled for internal terrace housing means that the total area of insulation required is significantly reduced, keeping costs down while representing a majority of the heat loss. Weight is added to this hypothesis by the under-representation of fully detached houses, which made up 13.6% of the stock but only 7.1% of the WHRSs with wall insulation, indicating that the additional surface area of an insulation-requiring wall reduces the probability of an optimal wall insulation retrofit.

The cumulative distribution of optimal wall insulation thicknesses can be found in Figure 5.16. The distribution of wall thicknesses is more even than that of roof installations, suggesting the material costs were not dominated by labour costs alone. The Pearson correlation coefficient (ρ) was used to inspect the attributes associated with insulation thickness among the subset of WHRSs of this kind. The same finding was observed here: buildings with a larger total wall area to be insulated were correlated with thinner

[2]Details of how lifecycle emissions savings are calculated are presented in Chapter 6.
[3]Excluding cavity wall insulation, which was present in 48.6% of the stock.

retrofit installations ($\rho = -0.32$). The glazing ratio of a building was also negatively correlated with optimal wall insulation thickness ($\rho = -0.29$), indicating that heavily glazed buildings are suited to thinner wall insulation. This is a fairly intuitive finding, as higher glazing ratios indicate a smaller proportion of heat is lost through the walls.

The second most common measure found in WHRSs was roof insulation, with 9.61% of the sample stock and 25.6% of the WHRSs including some degree of roof insulation. These were common co-retrofits, with 49.58% of these measures being part of a multi-measure WHRS. A key difference between wall and roof retrofits, in this case, is the significantly high prevalence of roof insulation within the existing stock. Indeed, the majority (79%) of the existing stock have more than the minimum thickness of roof insulation already installed, making the retrofit quite costly for a measure that only increases the thickness of the insulation of an existing area. Indeed, all of the roof insulation measures discovered were applied to buildings that either had no existing insulation or which had the minimum (30mm). While only this subset of the stock benefited from an installation, these WHRSs did obtain a mean NPV of £3,800 from an investment of £6,236 to achieve an annualised ROI of 2.44%. This coincided with a mean energy saving of 5006.52 kWh/year and a lifecycle emissions reduction of 29.01 Tons of $C0_2e$.

The cumulative distribution of optimal roof insulation thicknesses can be found in Figure 5.16. Unlike the wall insulation thicknesses, which were observed to be more evenly distributed, the roof thicknesses were clustered around the extremes. In 50% of the solutions, it was found that a minimal installation of 75mm would be optimal and in 19% of the WHRSs, it was found that the maximum thickness of 400mm would be optimal, with the remaining 31% distributed between these two values. One possible cause for these polarised optimal solutions is the nature of the cost model combined with the attributes of the houses being retrofitted. The majority of roof retrofits were performed on buildings with no existing roof insulation, with the remainder on retrofits with minimal thickness installations. When no roof insulation is present, a significant proportion of heat loss occurs through the roof and so even the installation of minimal roof insulation provides a significant benefit. However, there are diminishing returns to additional thickness on a roof with existing insulation, evidenced by the under-representation of these buildings in the WHRSs sample. Therefore, some dwellings will benefit from any kind of insulation, while those with existing insulation will only benefit from significant thickness increases. This hypothesis is borne out in the data, as all of the retrofits on dwellings with existing roof insulation increased their insulation to the maximum level. Naturally, all minimal installations are buildings with no existing roof insulation.

Heating fuel source retrofits had the highest NPV return of any retrofit when installed,

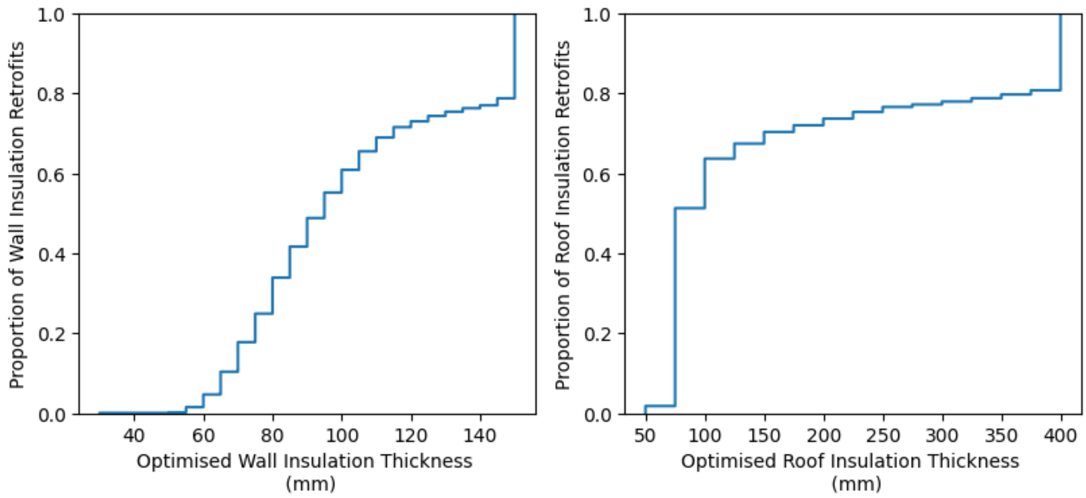

Figure 5.16: Cumulative distribution function of insulation-optimal wall and roof thicknesses.

with a transition from electric to gas heating resulting in an average NPV return of £42,000, which represents a significant ROI of 38%. However, these retrofits were only optimal in 2.29% of dwellings, given these measures were only available to electrically heated houses which made up the minority of the initial stock. It is also generally true that larger and higher energy use dwellings are less likely to be electrically heated, so the subset of dwellings that are large enough to recoup the investment in gas heating systems are likely to already be heated using that fuel source. As such, only 6.1% of the WHRSs included a heating system. None of the solutions involved retrofitting electric heating over gas central heating, which is unsurprising in the NPV-maximising scenario given the relative cost of electric heating is significantly higher than gas. The majority (84.72%) of heating installations were single-measure retrofits, with no other measures included. Due to the estimated decarbonisation of the grid during the lifecycle of the installations, these WHRSs were the only set to result in increased net emissions, averaging 10.37 tCO_2e of additional emissions.

Glazing was the least common retrofit in the sample, with only 6.52% of solutions including a glazing installation and 17.35% of the WHRSs including a glazing retrofit. Interestingly, 59.75% of glazing retrofits occurred alongside other retrofits instead of alone, making glazing the most commonly co-retrofit installation. This is likely due to the relatively low cost of installation compared with other options, which can be seen by the clustering of glazing retrofits at the lower end of the investment cost axis in Figure 5.17. The mean WHRS investment including glazing retrofits was £6,990 and resulted in mean energy savings of 4,930 kWh/year and a mean reduction in lifecycle emissions of 28.5 tCO_2e. These WHRSs returned an NPV of £3,270, resulting in an annualised return on investment of approximately 1.87%. When investigating the types of dwellings that found optimal glazing retrofits, the majority (86%) were initially single glazed with the remainder being double glazed dwellings. This is unsurprising, as single glazed properties receive the greatest benefit from additional glazing thickness. The population of WHRSs with optimal glazing installation also had a higher than baseline glazed area (30.46m^2 vs 27.33m^2) but a smaller average number of windows (5.30 vs 5.93). This indicates that dwellings with larger average window sizes were more likely to be good candidates for glazing installations. This is in line with expectations, as a labour and materials cost model penalises additional installations compared with installing more windows with a larger area. This is also the likely cause of the retrofit choices, with the vast majority of glazing retrofits being the installation of triple glazing as seen in Figure 5.17. Given the relatively small marginal cost of triple glazing compared to double glazing, solutions in which a glazing retrofit was viable tended to find the energy savings of the additional thickness exceeded the additional cost.

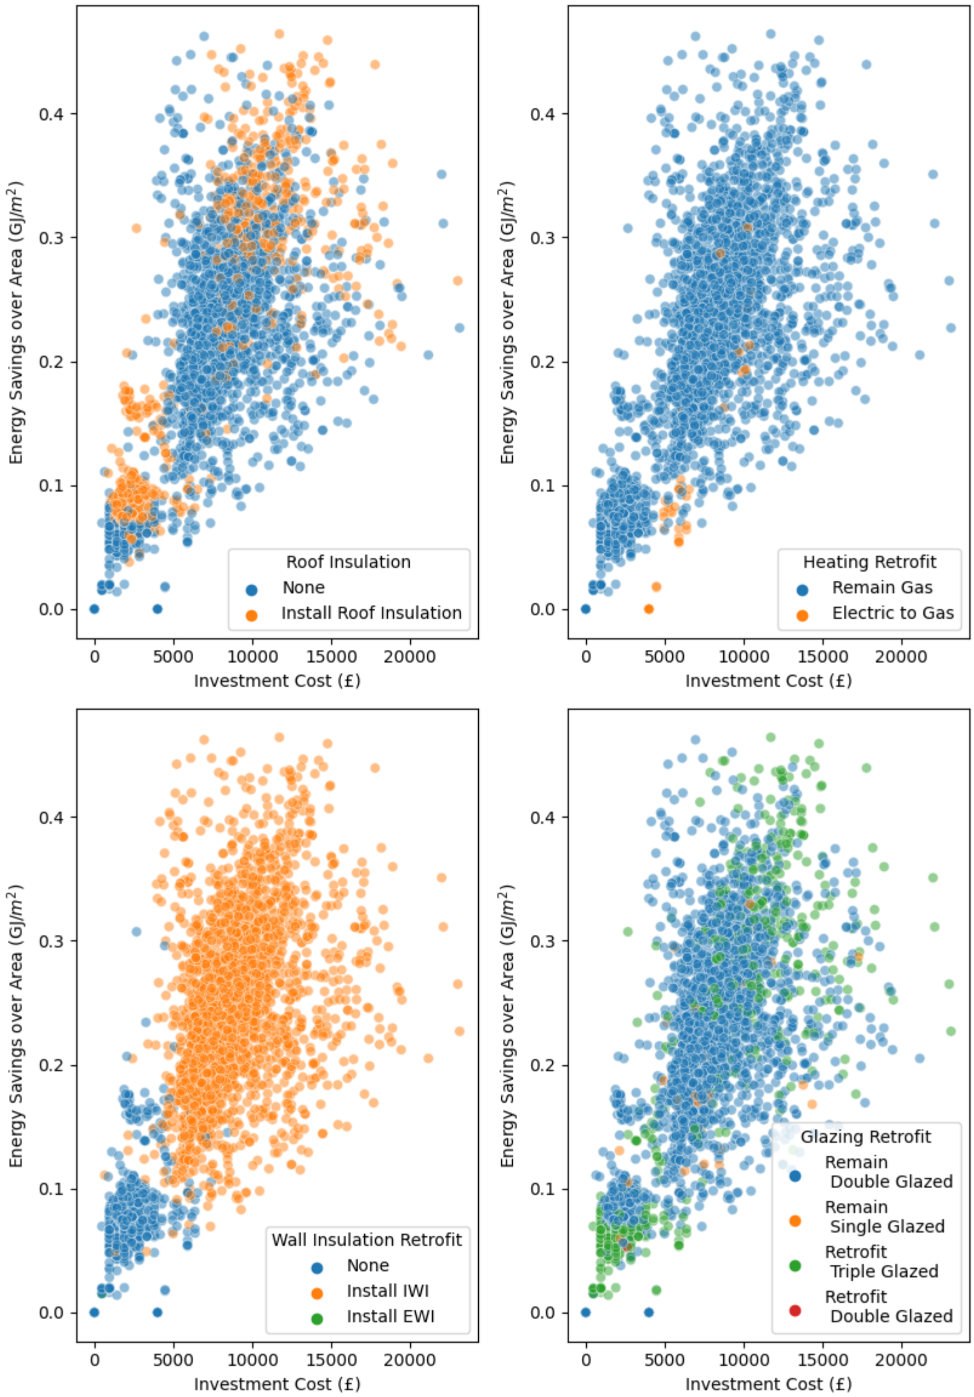

Figure 5.17: Scatter plots of the investment costs against energy savings per m^2 of floor space, coloured by the retrofit installation performed.

5.4 Single-Objective Surrogate Optimisation

The optimisation output generates a near-optimal retrofit per building optimised; this optimised state is highly dependent on the input building's initial state. To create a predictive Surrogate Optimiser (SO) (shown in Figure 5.18) a set of ANNs were trained. The sample of 10,000 profitable retrofits discussed in Section 5.3.3 represents the data set of solutions available for SO training, test and validation.

Some of the metrics used for model evaluation in this section have been used previously for the regression problem laid out in Section 5.2, such as the coefficient of determination (r^2), Root Means Square Error (RMSE) and Mean Absolute Error (MAE). However, this is the first classification problem introduced and so additional evaluation metrics are required to summarise the performance of classification models. Accuracy represents the number of correct predictions divided by the total number of predictions made. Precision represents the number of correct predictions over the total number of predictions for a given class. Poor precision indicated a large number of members of other classes are being mislabelled to the class at hand. Recall represents the percentage of correct predictions of a class over the number of actual members of that class in the data set. Poor recall on a class suggests that many true members of that class were mislabelled by the model. The final metric used in this section, the F1 score, is the harmonic mean of precision and recall: The product of precision and recall over the sum of precision and recall. This gives a metric that punishes models in which either precision or recall are poor. The advantage of using F1 score alongside accuracy is particularly stark when using unbalanced classes, as these models could achieve reasonable accuracy by predicting the modal class by default [195]. The F1 score will identify models with this issue more starkly than accuracy measurements. When evaluating multi-class problems these metrics can be presented in a weighted average format which calculates the metric for each class label, then provides a weighted average based on the frequency of that label in the test set.

5.4.1 Model Hyperparameter Tuning

The SO training varies from the SEPM described in Section 5.2.4 in a few key ways. Firstly, instead of training a single regression model, a set of different models need to be trained through both classification and regression due to the disparate data types of the target values. Multi-class classification models are required to estimate the desired material types while regression models are better suited to the continuous insulation thickness values under consideration. Within the models themselves, the types of pattern the model is attempting to learn will also vary. This was identified in Section 5.3.3, where it was observed that the characteristics of buildings with different measures differed considerably. Another difference between SEPM training and SO training is that these models are significantly newer, and as such, there is limited literature to provide sensible defaults or parameter ranges to explore.

Given the challenges provided by the hyperparameter tuning of the SO, an automated tuning mechanism was applied to the SO training workflow. The Keras-Tuner module was used to systematically sweep the hyperparameter space for the purpose of model training [196]. This allowed the hyperparameters of each model to be tuned without the requirement for an exhaustive grid search or manual intervention each time the models change. This second feature, removing the need for manual oversight, is likely to be particularly important as any changes to the underlying optimisation parameters may significantly impact the nature of a given problem. The hyperparameter tuning algorithm used for hyperparameter optimisation was Hyperband, a hyperparameter tuning algorithm that simultaneously trains a population of models and terminates poorly performing candidates early to allocate resources to high-performing candidates [197]. This algorithm was chosen for the resource efficiency of the tuning process which, given the conceptualisation of SO as a modelling workflow coming with the necessity to retrain and re-tune the model without the requirement of domain knowledge, is an important property.

The list of hyperparameter settings explored can be seen in Table 5.8. As with the SEPM training discussed in Section 5.2.4, the pre-processing procedures and training batch sizes were added to the parameter search space. Additionally, a higher degree of freedom was applied to the layer parameters, allowing 1-7 layers of sizes between 8 and 256 neurons to be selected by the hyperparameter tuning process. This higher freedom in layer selection is necessary as there are significantly fewer published models in this application compared with SEPMs and as such, there is less certainty about which network architectures would be considered sensible defaults to attempt. The Keras-Tuner implementation of Hyperband was implemented, with a max epochs parameter of 100 and a branching factor of 2, resulting in 359 trials per model [196]. Validation loss, as specified in Table 5.8 was used for the tuner's optimisation objective and a validation set of 2,000 buildings was used. An early stopping procedure stalled a given model's training if the validation loss stalled below 0.0001 over a patience of 10 epochs.

The results of the parameter tuning are specific to the training set and scenario presented and varied fairly significantly between the models. One common factor across all of the trained hyperparameters was the faster learning rate of 0.01, the same as the 'sensible default' used for SEPM training. Given that this result was found to be robust across multiple re-runs of the parameter tuning procedure, it can safely be set to a default value to reduce the search space during re-runs. The other hyperparameters varied between the models, with the most significant variance being the network structure. As may be expected, the problems that we would anticipate to be the easiest from the discussion in Section 5.3.3, generally resulted in a smaller number of layers which were themselves smaller. In contrast, the more challenging problems such as the roof thickness regression model benefited from the full 5 layers (each layer exceeding 200 neurons), the largest of the networks.

Setting	Regression Models	Classification Models
Loss Function	L2 (Least Square Errors)	Categorical Crossentropy
Last Layer Activation	Linear	Softmax
Pre-proccessing	Normalisation, Logarithmic, Standardisation	
Number of Layers	1 - 7 (Steps of 1)	
Neurons Per Layer	8-256 (Steps of 8)	
Activation Function	ReLU	
Learning Rate	0.01, 0.001, 0.0001	
Training Function	Adam (alpha=0.001, beta1=0.9, beta2=0.999 and epsilon=10E-8)	

Table 5.8: Description of SO hyperparameter ranges used for tuning.

5.4.2 Sample Size

In order to test the effect of the training set sample size on model performance, models were trained using a sample size varying from 500 to the full 10,000 set. Models were trained using their corresponding tuned hyperparameters, determined from the tuning stage. A total of 20 repetitions were performed for each sample size to account for the stochastic model training procedure. Given the 8 measures which comprise a potential WHRS, a total of 3,040 models were trained.

The effect of sample size on trained model performance can be seen in Figure 5.19. It should be noted that the performance metric evaluated for the classification models, weighted f-1 score, increases with model performance while the RMSE used for regression performance decreases as the model improves. It can be seen that all models, with the possible exception of the EWI material classification, improve as a greater sample size of near-optimal WHRSs is used to train them. EWI material classification is likely an exception to this due to the small number of samples found in the training data set, as discussed in Section 5.3.3, which renders model performance evaluation subject to more stochasticity due to its sparsity in the training and subsequent test set. The other models noticeably increased in performance with additional sample size, although the diminishing returns that were experienced, as seen in the curved y log(x), best fit lines plotted over the results.

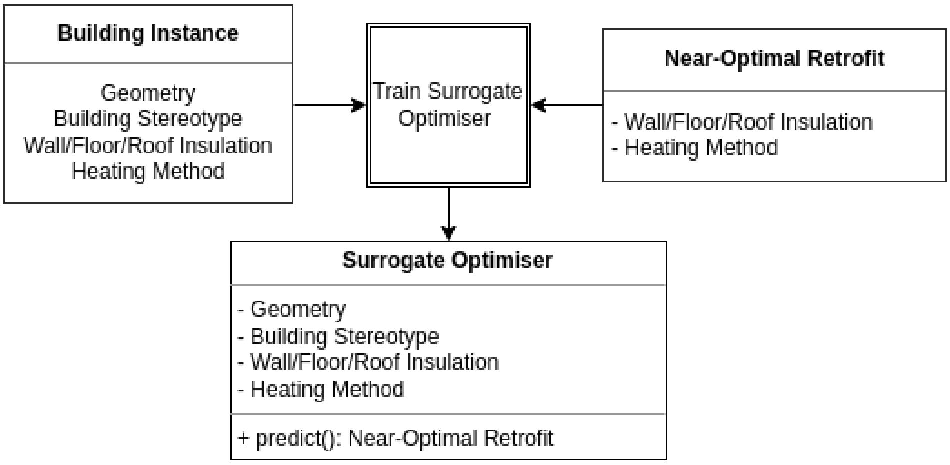

Figure 5.18: Process flow for training the Surrogate Optimiser.

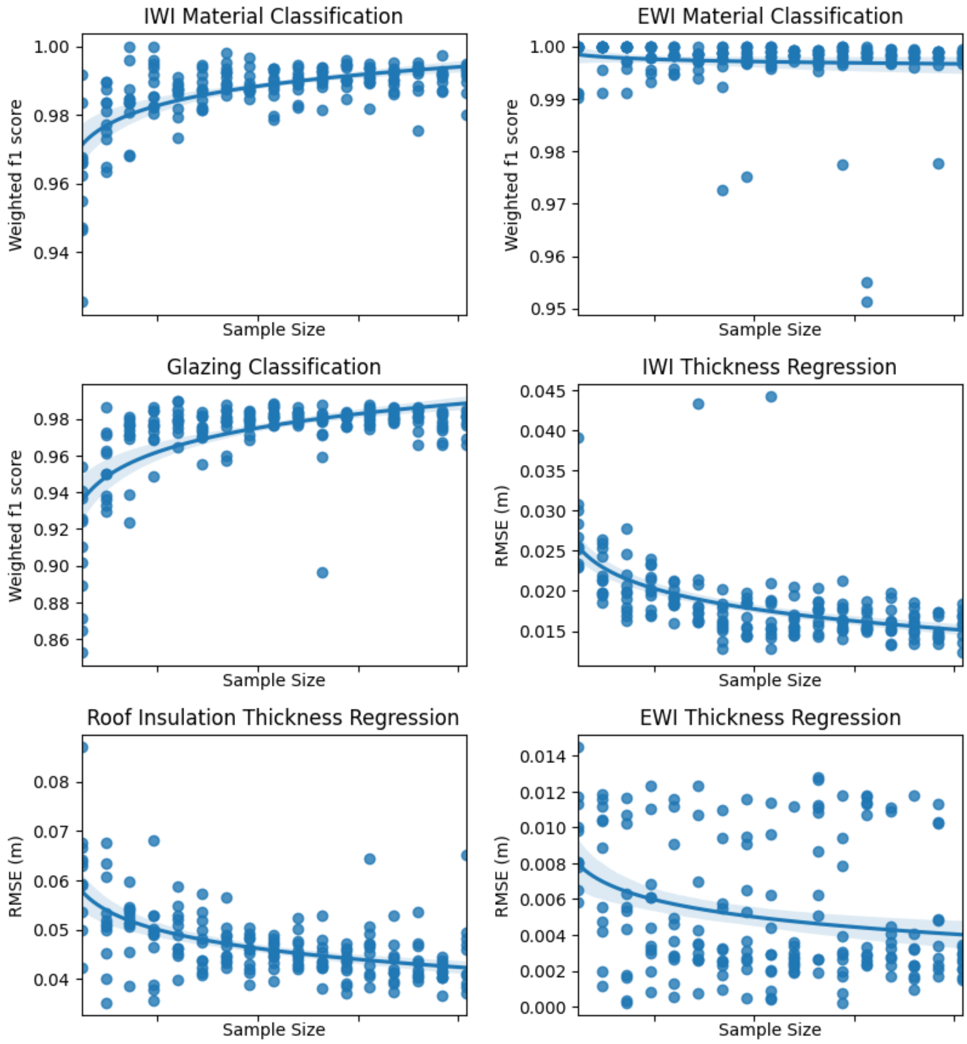

Figure 5.19: Effect of sample size on SO performance, showing diminishing but positive improvements between 500 and 10,000 training sets.

5.4.3 Surrogate Optimiser Results

Glazing retrofits were the least common installation, as discussed in Section 5.3.3. This provided the smallest training sample with only 600 WHRSs in the 10,000 buildings sample, leaving 9,400 in which no retrofit was performed. Additionally, this measure was most commonly a co-retrofit, with approximately 60% of these retrofits occurring alongside another measure. This makes false-positive retrofit classifications more difficult to

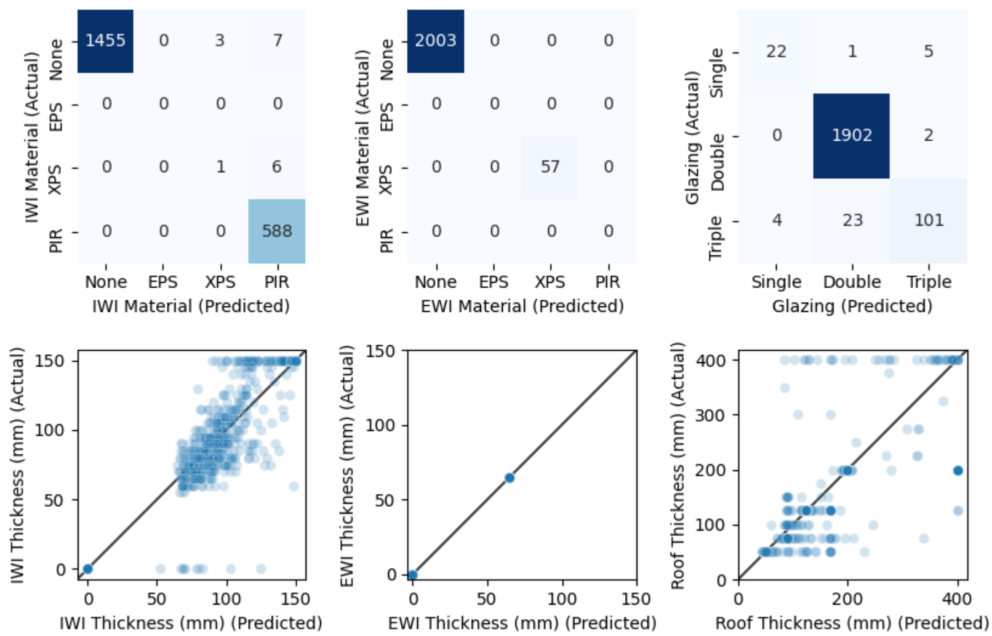

Figure 5.20: Confusion matrices (top) and Actual vs Predicted scatter plots (bottom) of Surrogate Optimiser training set (N = 2060).

spot, as the WHRS may still achieve a positive NPV even in instances where the glazing retrofit is not beneficial. Despite these challenges, glazing classification performance was good. The confusion matrix for glazing classification is shown in the top right of Figure 5.20, with the gene values of 0, 1 and 2 representing single, double and triple glazing values. The double glazing class, most of which are preexisting and therefore represent no retrofit, performed very well, with an F1 score of 0.99, indicating that most of this class was correctly classified. The single and triple glazing classes are of most interest, however, as these represented the bulk of retrofit installations found. Single glazing obtained a precision of 0.89 and recall of 0.78 while triple glazing obtained precision and recall scores of 0.89 and 0.83 respectively. These values indicate that approximately 17% of profitable triple glazing retrofits are being missed while 22% of single glaze properties, for which a glazing retrofit is not profitable, are being inappropriately designated a retrofit. Given that 40% of glazing retrofits are single retrofit measures, for which misapplied retrofits can be cheaply identified, this leaves 13.2% of dwellings that are optimally single glazed being sub-optimally retrofitted. However, since this represents only 2% of the sample, even across the entire city data set of 95,500 dwellings the total expected number of errors would be around 250.

Roof insulation retrofit predictions were made somewhat challenging by the clustering of thicknesses at distinct but key values as opposed to the more continuous distribution of IWI insulation thicknesses. This was discussed in Section 5.3.3 and was visualised in Figure 5.16. This poses a challenge, as a regression model in which many results are polarised at the extremes of the possible range may result in more drastic errors. This is observed to some extent in the results of the roof insulation thickness model, shown at the bottom right of Figure 5.20 by the horizontal line across the top of the scatter plot. This line of points represents maximum actual optimised thicknesses that the model has incorrectly predicted to be sub-maximal. Despite some errors, particularly at the extremes of the model, the overall performance of the regression model is good. An adjusted r^2 score of 0.82 indicates that the majority of the variance of the predicted insulation thickness is explained by the model inputs. The residuals, which are shown in Figure 5.21, appear to be normally distributed and centred around 0. This distribution hypothesis was not rejected, with a Kolmogorov–Smirnov test for normality at a 5% significance level. The MAE of the predictions is 8.3mm, although the RMSE is 24.1mm due to the disproportionate size of some of the errors made. These values are encouraging, as the values are to be coerced back into discrete thicknesses when considering which retrofit measure to install, with thicknesses differing by 25mm which is larger than both the MAE and the RMSE. When these regression values are converted back to the thickness gene values used for the optimisation stage, it is possible to evaluate the model as a classification procedure. When considering these classes, the roof thickness model scores an accuracy of 0.88 and a weighted F1 score of 0.91, indicating good classification performance across classes. Where the thickness is misclassified, it was often by only a single thickness class. A confusion matrix of roof thicknesses coerced back into classes can be seen in Figure A.1 of Appendix A.3.

Wall insulation retrofits were classified separately for EWI and IWI materials, although as stated in Section 5.3.3 the majority of installations in this scenario were IWI, with only a small minority of WHRSs containing an EWI, most of which are preexisting in the stock. This can be seen in the confusion matrix and thickness plot at the centre of Figure 5.20, which demonstrates that a small number of optimal EWI installations are identified well by the model, and the thicknesses are also predicted accurately. As EWI installations will become more prevalent in Chapter 6, the model has been retained for completeness. Should the modelling parameters used in Section 5.3 change, this essentially dormant model would again be required.

In contrast to the sparse EWI installations, IWI installations were the most abundant in the optimisation data set. The most common installation material in this scenario was polyisocyanurate (PIR), represented by class 3 in the top left confusion matrix of Figure 5.20. The SO was accurately able to identify almost all of the solutions which included a WHRS, although due to the heavily unbalanced classes the small number of WHRSs which required alternative materials were incorrectly classified as PIR installations. It is possible, however, that the small number of these installations indicate an error at the previous stage (optimisation). It is also worth noting that these misclassified IWI material retrofits were always classified to PIR installations as opposed to no installation at all (signified by class 0 in the confusion matrix), showing the model was able to accurately identify all IWI retrofits even when the material was classified incorrectly.

The thickness of optimal IWI installations was predicted using a separate regression model. Unlike the roof insulation thickness, which was clustered at specific points, the optimal wall insulation thicknesses were more evenly distributed across the possible range as discussed in Section 5.3.3 and visualised in Figure 5.16. The regression model performed well, achieving an r^2 of 0.951, a MAE of 4.03mm, and a RMSE of 8.85mm. Figure 5.20 shows a scatter plot of actual optimised thicknesses plotted against predictions, with most points sitting close to the line of equality. The residuals of the model are shown in Figure 5.21 and appear to be normally distributed and centred around 0. This distribution hypothesis was not rejected with a Kolmogorov–Smirnov test for normality at a 5% significance level. While the regression results are promising, the discrete IWI thicknesses differ by just 5mm, and so when values are coerced back into classes some of the predictions are classified incorrectly. The model's classification accuracy of 0.78 is good, however, and this is paired with a weighted average F1 score of 0.76, indicating good levels of recall and precision across classes. Where thickness classes are incorrect, the error is typically within one ordinal class, indicating the magnitude of the error is only small. A confusion matrix of IWI thicknesses coerced back into classes can be seen in Figure A.1 of Appendix A.3.

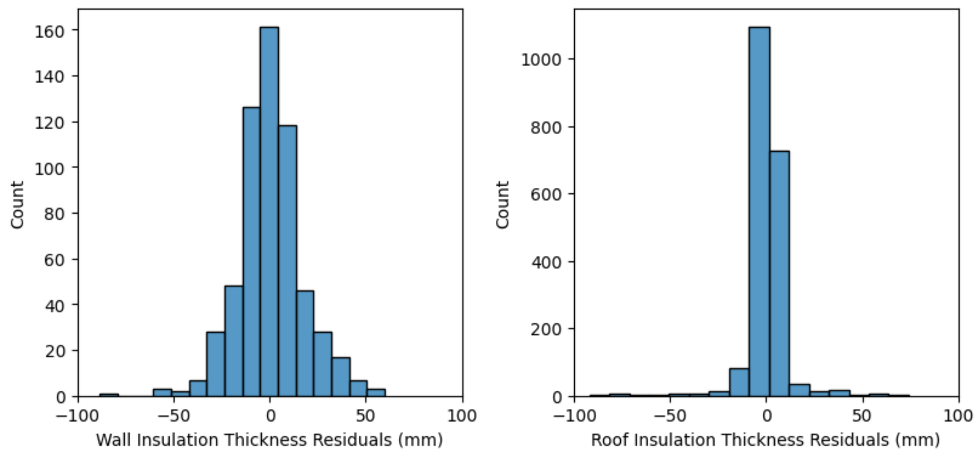

Figure 5.21: Histogram of wall and roof insulation thickness residuals from Surrogate Optimiser predictions.

5.5 Application of Single-Objective Surrogate Optimiser

The Surrogate Optimiser (SO) performance, evaluated in Section 5.4.3, showed that the input features carried sufficient data to make Whole House Retrofit Solutions (WHRSs) with an acceptable level of accuracy. The next step is to apply the SO to untrained data to analyse the proportion of the stock that was not used for training and static testing. The SO will be applied in two ways. In Section 5.5.1, a static analysis of retrofit potential will be performed by applying the SO to the entire building data set and analysing the results. In Section 5.5.2, the SO will be integrated into a simple Agent-Based Model (ABM), which will be introduced with the intention to expand its functionality to environmentally conscious agents in Chapter 6.

5.5.1 Static Whole Stock Analysis

A static evaluation of the building stock involves predicting the near-optimal WHRSs for all buildings in the data set. The predicted solutions can then be evaluated using the same NPV model that was used in Section 5.3 to create the initial sample of near-optimal solutions. This allows a comparison between the sample of traditionally optimised WHRSs and those predicted by the SO on the unknown solutions. This can be considered analogous to the upscaling procedure common in archetype-based housing stock models discussed in Section 3.2 of the related works. However, given the procedure uses the discovered features of the ANN for retrofit predictions, there may be differences between the traditionally optimised sample and the predictions of the whole population. In this sense, the whole stock evaluation procedure is an additional process of performance testing for the SO as well as an evaluation of the potential within the whole stock.

The procedure for evaluation of WHRS-performance was the same as that used in the GA. Initial energy demand was calculated using the SEPM introduced in Section 5.2. The SO was then used to predict WHRSs for the entire building stock based on each dwelling's features. These retrofits were then applied to the dwellings and new energy demand was calculated using the SEPM again. The cost and savings models were applied based on the retrofit solution and energy savings for each dwelling, then the NPV was calculated using the discount rate. In total, this requires 3 calls to DNN models per building and a single call to the financial models for each dwelling in the data set, a significant computational saving over the original optimisation procedure. Indeed, the total wall time required to perform this evaluation on the 95,500 dwellings stock is approximately 30 seconds, including loading the required data into memory.

5.5.1.a Negative NPV Solutions

Prior to discussing the quality of the full set of solutions, it is worth discussing a particular subset of predicted solutions: those for which the NPV was negative. Given that not performing a retrofit was always a valid decision at the GA optimisation stage, none of the optimal solutions had negative NPVs. Not performing an installation has an intrinsic NPV of zero which therefore always dominates negatively performing solutions among this breed of self-interested rational agents. However, 1.79% of the sample were predicted as negatively scoring retrofit solutions. Given these retrofits can be scored in a computationally cheap manner, these misapplied retrofits are simple to identify and remove. However, there is potential to miss some true retrofits by discarding these solutions, as it is possible that a positive NPV solution exists but was not correctly identified. The mean NPV of these failed solutions was £-1,305.70 from an average investment of £5,842.75. The average energy saving of these retrofits was only 15.6 kWh/m^2 compared with 54.5 kWh/m^2 of the positive NPV WHRSs, indicating that these retrofits performed significantly worse than expected by the model. This is also reflected in energy savings over investment cost, with failed retrofits obtaining energy savings of 0.435 kWh per pound while positively performing installations saved 0.70 kWh per pound of investment.

$$V_{u0} = \frac{-NVP}{LCS} \tag{5.3}$$

Equation 5.3: The minimum indifference carbon value is the negative ratio of NPV and LCS.

While all of the solutions with negative NPV would be rejected by the self-interested agent decision model used in this chapter, it is worth considering how environmentally conscious agents may consider them. The scatter plot in Figure 5.22 shows the NPV and net carbon emissions savings of these rejected retrofit predictions. While it can be seen that some predictions perform extremely poorly, with both low NPV and little environmental impact, there are a significant number of solutions on the right-hand side of the graph. These have a small negative NPV but significant lifecycle CO2e savings. While we SO is limited at present to environmentally indifferent agents, it is possible to consider how environmentally conscious decision makers would need to be to consider these solutions. To calculate this, we can use the minimum indifference carbon value: the minimum financial value the agent must place on each tCO2e for them to be indifferent between that solution and inaction. The formula for minimum indifference carbon value is shown in Equation 5.3. The distribution of the minimum indifference carbon values of these solutions is shown in Figure 5.22 and demonstrates that many of these solutions, which are not valid with rational agents, would become feasible with agents who have even a small environmental preference.

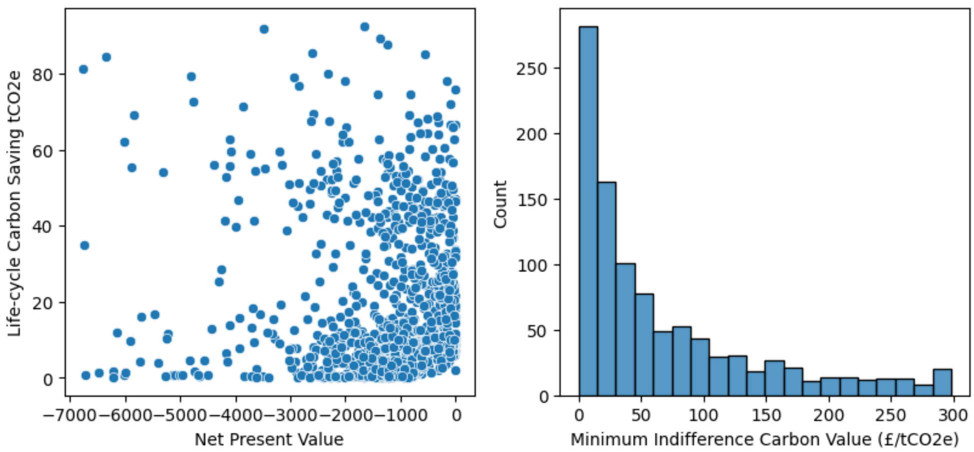

Figure 5.22: Scatter plot of NPV and LCS of negatively performing retrofit predictions (left) and a histogram of their corresponding minimum indifference carbon values (right).

5.5.1.b Positive NPV Solutions

	GA Optimised Sample	SO Predictions
Sample Size	10,000	95,500
Percentage of Dwellings with WHRSs	37.40%	37.02%
Percentage of Co-retrofits	20.72%	20.34%
Average WHRS Investment Cost	£7,112.00	£7,259.00
Average WHRS Net Present Value	£5,148.00	£4,797.00
Average Carbon Saving (tCO2e)	28.31	28.85
Average ROI	2.86%	2.64%
Dwellings with Glazing Retrofits	6.52%	5.54%
Dwellings with Wall Insulation Retrofits	27.50%	28.02%
Dwellings with Roof Insulation Retrofits	9.61%	9.58%

Table 5.9: Comparison between a Genetic Algorithm optimised sample and the Surrogate Optimiser predictions of whole stock.

After considering the minority solutions erroneously predicted by the SO, the positively scoring solutions can be examined. At this stage in the analysis, the solutions identified above were removed, with the dwellings added to the subset of properties for which no positively scoring solutions were discovered. The resulting data set of positive solutions can now be analysed and, importantly, compared with the sample of traditionally optimised WHRSs discussed in Section 5.3.3. Table 5.9 compares some key statistics from the optimised sample discovered by the GA, alongside the batch of predicted solutions created by the SO. The comparison of results shows that for most elements the SO predictions mirrored the GA sample quite closely. The total proportion of WHRSs representing 37.4% of the GA sample and 37% of the SO predictions, indicating that the vast majority of WHRSs present in the underlying stock have been identified. There is some difference in the nature of the solutions discovered though, with the SO sample requiring a slightly higher investment of £7,259 compared with £7,112 required by the GA, which also corresponds to a reduced average NPV (£4,797 to £5,148). This slight tendency of SO predictions to over-invest did manifest an increase in lifecycle emissions savings (28.85tCO2e to 28.31tCO2e) which indicates that the over-investment was in environmentally-friendly measures and also corresponded to net energy savings. When considering the breakdown of different retrofit types, the number of WHRSs with roof insulation mirrored the GA sample very closely, while IWI measures were slightly over-represented and glazing installations slightly under-predicted (See Table 5.9). The SO was also efficient at identifying co-retrofit solutions, with a similar proportion predicted in the whole stock as were found in the GA sample (20.72% to 20.34%). It should be considered, however, that the co-retrofit WHRSs represent a potential area of hidden errors, as an effective retrofit paired with an inefficient one is harder to identify compared with an inefficient retrofit on its own, such as those discussed in Section 5.5.1.a.

The thicknesses of the predicted retrofits can be seen in Figure 5.23. These distributions mirror those identified by the GA sample which was shown in Figure 5.16. As with the prior distributions, the roof insulation thicknesses cluster around the upper and lower ends while the wall insulation thicknesses are more evenly distributed across the available thickness range. The main difference in distribution is the slight reduction in maximal thickness roof installations, an expected result given the misclassification of some of these thickness classes observed in the SO training data discussed in Section 5.4.3. However, these errors were often by only a small number of classes, and this can be observed in the corresponding increase in the final two thicknesses classes.

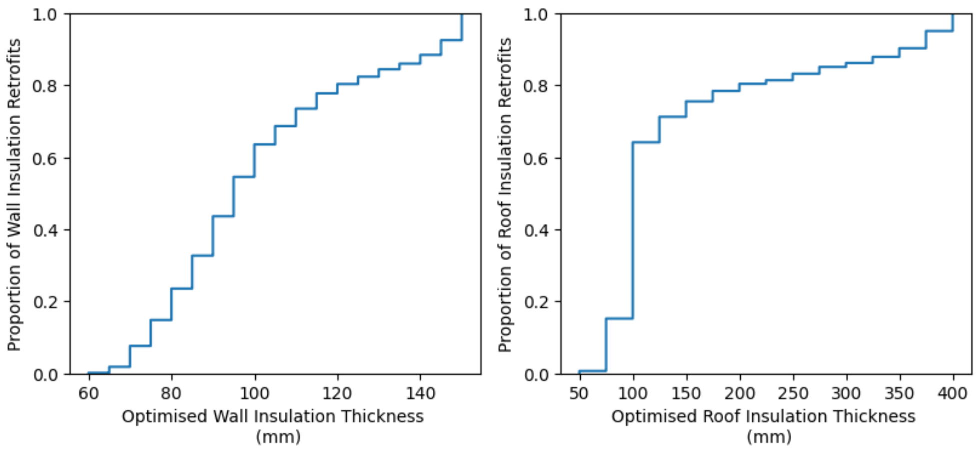

Figure 5.23: Cumulative distribution of predicted optimal insulation thicknesses.

5.5.2 Dynamic Analysis with an Agent-Based Model

While the static analysis of retrofit potential in the entire city data set is of interest, similar aggregate values can be obtained using more traditional sampling and upscaling techniques. One of the key benefits of a SO approach is the ability to predict a retrofit decision for specific households in real-time. To demonstrate this, the SO trained in Section 5.4 will be integrated into a very simple Agent-Based Model (ABM) to observe how the stock may adapt over time.

The model introduced in this section will be a simplistic one which will be expanded in Chapter 6 to capture more complex behaviours relating to an agent's individual carbon preferences. In a sense, the agents at this stage are heterogeneous: They are all attempting to maximise their NPV in considering a retrofit. However, each agent is faced with a different problem as all of the dwellings in the building data set are unique. Performing retrofit evaluations at the scale of an entire city using the optimisation procedures performed in Section 5.3 would be infeasible, as the computational time required to optimise and evaluate each dwelling would be significant.

5.5.2.a ABM Design

The ABM design at this stage is kept very simple. The main parameter introduced is the evaluation rate, which represents the probability that a given agent will evaluate their dwelling's retrofit potential within a given year. Agents who do decide to evaluate their dwelling for retrofit have a WHRS predicted by the SO, which is then scored and then installed if the NPV is positive and rejected if the retrofit does not provide a positive lifetime value. This simple agent-decision model is visualised in Figure 5.24.

Three scenarios will be performed with different evaluation rates. These have been conceptualised in terms of outreach campaigns. In the baseline scenario, an evaluation rate of 1% will be used and will be compared to a standard outreach program that achieves an evaluation rate of 1.5%, and a significant outreach program that achieves an evaluation rate of 2%. These programs can be conceptualised in terms of marketing campaigns by government departments, retrofit installers, or other events that cause increased desire for households to evaluate their dwelling's retrofit potential.

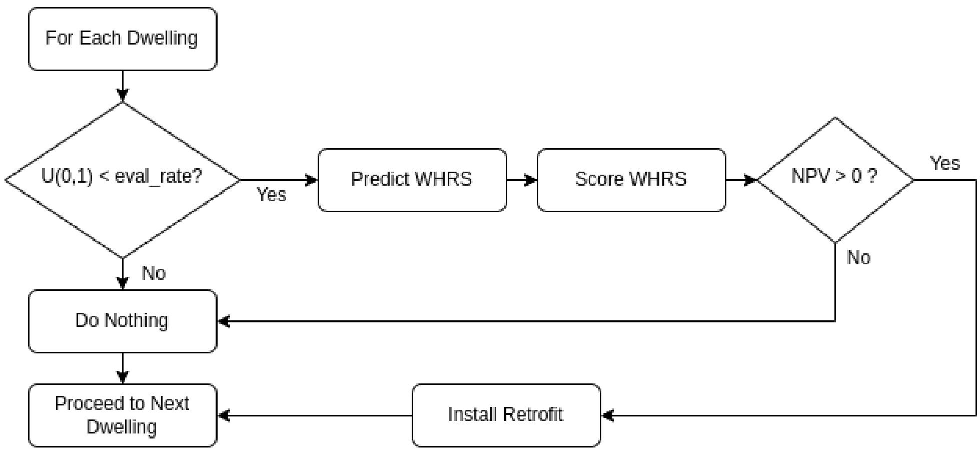

Figure 5.24: Flowchart with a high level description of the decision process each dwelling undergoes during each time period. The evaluation rate is represented by a float called eval_rate.

5.5.2.b ABM Results

The ABM was run for a period of 25 years, with each scenario repeated 100 times in order to gauge the level of output variance. The results of the model are visualised as a series of time-series plots of key aggregate outcomes in Figure 5.25. As would be expected from the design of the model and the relative lack of dynamic components, the outcomes are generally proportional to the relative evaluation rate of the scenario at hand. There are some points of interest, however. The scatter plots, which plot each year of each repetition of the scenarios, show the level of variance in the underlying selections, with the NPV showing significantly more variance than the investment costs within each year. The variance of the NPV is such that the best performing years of the baseline scenario outperform the worst years of the outreach scenarios. Nonetheless, the total number of retrofits is significantly higher when outreach programs are in place. While this model is largely demonstrative, there are some results to consider. Consider, for example, how much should be invested into each outreach campaign by a government that was interested in stimulating the economy through increased infrastructure investment. The average annual investment in the baseline scenario is £2.7 million compared with £4.0 million and £5.4 million in the outreach and significant outreach scenarios respectively. Knowing that each level of outreach would return a £1.3 million increase in investment, as well as knowing the distribution of investment spending from the stochastic runs of the model, could aid in policy decision making when planning these scenarios.

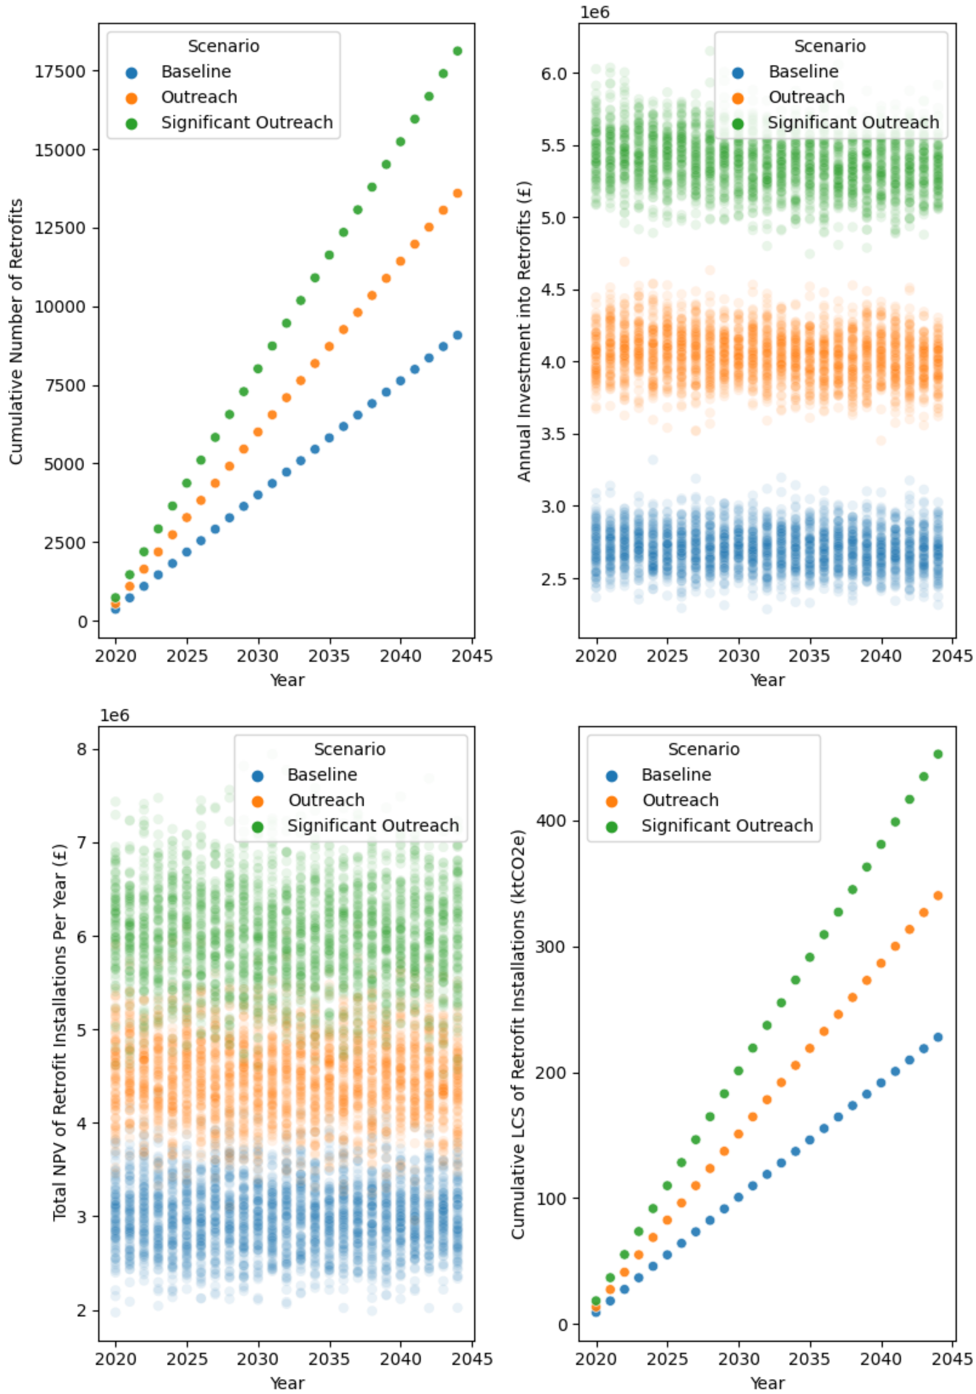

Figure 5.25: Time series plots of 300 ABM simulations showing the effect of retrofit outreach programs are significant. Cumulative plots show the annualised mean.

5.6 Discussion

The following section highlights and expands upon some points from the research thus far that warrant further discussion. A particular focus will be placed on those points of discussion that are specific to the problem of self-interested and purely rational decision makers outlined in this chapter, with general discussions on the SO technique relegated to the general discussion in Chapter 8.

5.6.1 The benefits of Surrogate Optimisation

A question could be posed at this stage regarding what the benefits of Surrogate Optimisation as presented in this chapter are over alternative urban-scale retrofit modelling techniques. The framing of the benefits depend upon the technique to which they are compared, and so they will be broken down and comparisons made.

Top-down modelling methods were discussed in Section 2.1.3.a and 3.2. Unlike bottom-up approaches, such as a SO which investigates individual sub-problems that make up a large system before aggregating them to measure whole-system behaviour, top-down models operate only on the aggregate behaviour. This can be done with the application of statistical or theoretical models that attempt to capture the aggregated behaviours without requiring or providing the granularity that bottom-up approaches utilise. These models are generally more computationally feasible, as they do not require individual evaluation of each dwelling, however, they have significant drawbacks over top-down approaches too. One key drawback is the lack of granularity in the results. The 95,500 buildings examined in this chapter are all unique and, after optimisation, the attributes of the properties and their retrofit potential are evaluated individually. This allows an examination of which dwellings or types of dwellings apply which retrofit, the nature of those retrofit measures, and the individual benefits those dwellings receive. This contrasts with much of the related works, which applied only a small and pre-selected set of solutions to all dwellings regardless of their properties. A second major advantage of the bottom-up SO approach is the level of control granted to modellers in controlling the simulation. Chapter 6 will take advantage of the bottom-up nature of the SO by instantiating households with individual preferences that affect their retrofit decisions, something that is not possible in top-down approaches which do not model individual agents and therefore cannot emulate their heterogeneity.

While top-down models have major disadvantages compared with the SO approach, we could compare the approach instead to what we shall term a pure bottom-up approach. This procedure involves applying the same or similar optimisation process as was used in Section 5.3 to individually discover the optimal retrofit for a given dwelling using traditional optimisation procedures. This method is feasible for small stocks, as was demonstrated by our collection of 10,000 sample retrofits for SO training. However, this approach is very computationally costly with each sample retrofit taking approximately 30 cpu seconds to obtain[4]. This equates to approximately 800 cpu hours to perform a single optimisation run on the entire stock. While this is not impossible, it is infeasible for procedures that require frequent re-running or evaluation, such as during a stochastic ABM. This problem is exacerbated when heterogeneous decision makers are introduced, as this prevents a calculate, store and retrieve approach that could be used when decisions are made from a single objective. The infeasibility of this pure optimisation approach increases with the size of the stock, as the simulation of the national building stock of the UK's approximately 24 million domestic dwellings would take 200,000 cpu hours for a single optimisation run [198].

While the pure bottom-up approach of simulating and optimising each dwelling grows infeasible as stock size and the number of repetitions required grows, there are alternative bottom-up approaches against which to compare the SO method. The most comparable of these is the process of sampling, archetyping, optimising and upscaling. A good example of this methodology is the EnHub platform [94], which uses a data mining approach to form archetypes for simulation before extrapolating those results to the entire stock. This method is significantly more computationally feasible than a pure bottom-up approach but does also face significant drawbacks. The responsibility of forming archetypes lies on the modeller, requiring significant research time as well as domain knowledge to form and sample properly. Also, while these approaches allow for some heterogeneity, they still rely on representative dwellings for given archetypes, reducing the granularity of the model and the diversity of outcomes. Consider the diversity of outcomes visualised in Figure 5.17 of Chapter 5.3.3 being forced into a small and discrete number of archetypes. This would also limit the heterogeneity of agents in any ABMs formed due to archetype constraints. It is worth discussing how these techniques compare to the SO methodology, which relies on the trained model to identify which underlying features of dwellings contribute to a given WHRS's optimality and for making predictions based on the patterns discovered. In this sense, the upscaling procedure performed can be considered a higher resolution and automated version of the upscaling that is manually performed with an archetype measure. The increased resolution and automation are both key advantages, removing the requirement of significant domain knowledge in the application as well as allowing individual modelling of households,

[4]On a workstation running an AMD Ryzen 9 3900X with 16GB of RAM.

granting more control. A key downside of the SO approach is the loss of interpretability compared with constructing archetypes manually, as the underlying data patterns are hidden within a black box model. This requires additional research to understand or interpret the models themselves.

5.6.2 The Assumption of Pure Rationality

Pure rationality, as we have defined it for the purposes of this chapter, means that household agents seek out WHRSs which optimise their personal financial returns, measured in this case by the NPV of the investment. Given the intention to relax this assumption in Chapter 6, the purpose of constructing these types of agents may be questioned, particularly as it can be observed that individuals do care about environmental concerns in an empirically measurable way [134, 145, 147]. One reason to consider purely rational agents was as a proof of concept for a smaller part of what will become a more complex problem. The restriction of the objective function keeps the scope of the model down and allows for an initial prototype to be examined in isolation. Another reason to consider purely rational agents is that it is an appropriate model for some of the population. While there are many environmentally conscious households, there are also decision makers for whom ROI is the key factor.

5.6.3 Model Assumptions and Parameters

When discussing the practice of modelling and simulation in Section 2.1.1, models were considered in terms of target systems with the behaviour of the model mirroring, in some way, the behaviour of the system they were targeting. A variety of different models were introduced in this chapter. The EnergyPlus simulation models were designed to capture the energy demand behaviour of specific buildings. The SEPM took those initial energy models as its target system, attempting to capture the model's behaviour in a less computationally expensive manner. The optimisation procedure introduced an objective function calculation of NPV that is in itself a financial model, attempting to capture the costs and benefits of a given installation under the circumstances presented. The optimisation process, as well as encapsulating the financial and energy models required to evaluate solutions, can be considered a model too as it attempts to capture real-world behaviour: a purely rational agent attempting to install the best retrofit for their objective.

It is clear that there are a significant number of models used in this procedure. In

Section 2.1.1 a mannequin was introduced as an example of a model, with the purpose of demonstrating how given clothes would look on a living human. This was our distinction between a model and a simulation, as it was only when the clothes were applied to the model that the purpose of a simulation is fulfilled, and the choice of clothes has a significant impact on the outcome. This is also true with all the models introduced in this chapter. Where models have been introduced, parameters and assumptions baked into the modelling have been explained but it is important to highlight that the results are valid only under those parameters and assumptions. This is also why this method is presented in terms of the method and model construction workflow, instead of the results of a given set of input parameters. Consider, for example, that a supply chain issue caused the price of a given construction material to increase significantly. This would impact the financial model that makes up the objective function of the optimisation stage, requiring retraining of the SO and revaluation of the potential within the stock.

Considering this method as a computation workflow, as opposed to the construction of a specific parameterised model, has implications at a few stages. Over-tuning the models for a given scenario's parameters is likely to be a waste of resources, as the hyperparameters and optimisation settings tuned are likely to need changing again when the scenario parameters are adjusted. The SO model selection presented in Section 5.4 could have included the removal of the rare or unused classes during model training to provide a small boost to model accuracy, however, this would come at the cost of generalisability. As will be seen in Chapter 6, when the parameters of the objective function model are changed, the classes that could have been removed here become useful or sometimes dominant. Another consideration of framing the method in terms of a workflow is the important fact that modellers or stakeholders using the workflow are aware of which parameter changes have the largest effects on computational cost. At the first stage, the EnergyPlus simulation models will have the largest impact in terms of model reconstruction, while parameters in the outer retrofit adoption model are the least costly. A high-level view of these stages and the retraining required to achieve them can be obtained by looking back to Figure 5.2.

5.6.4 Optimisation Problem Complexity

The complexity of the chosen optimisation problem is worth noting here. The optimisation problem itself was a combination of eight components varying from binary (in the case of heating) to 25 discrete thicknesses of IWI and EWI insulation. The total unconstrained optimisation search space was 1,800,000, making the problem significantly too substantial to solve at run time for each dwelling during a dynamic agent-based model of the type discussed in Section 5.5. It is worth noting, however, that for a given scenario setting, the range of feasible solutions does become significantly reduced due to the cost-benefit return of certain measures under any given scenario. As was seen in the scenario presented in this chapter and discussed in Section 5.3.3, the installation of EWI was generally (although not completely) dominated by IWI installations. This essentially removed EWI from major consideration when training the SO and therefore reduced the problem's effective complexity to 67,200 feasible solutions. It would have been possible to remove those attributes and perform the optimisation procedures again. However, the attributes were left in and the SO was trained including them in this section for a few reasons. Firstly, the SO method laid out in Chapter 4 and implemented in this chapter should be considered a computational workflow and the outcomes depend significantly on the choice of simulation parameters, such as the cost of labour and materials. It should be possible to adjust these elements without having to redesign the workflow itself, as doing so is labour intensive for a modeller as well as requiring potential domain knowledge. Automating the workflow also allows more extensive experiments to be carried out, as the stages can be re-run using data derived or altered programmatically at runtime without requiring human input.

Chapter 6

Modelling Environmentally Conscious Agents Using Multi-Objective Surrogate Optimisation

"There are no solutions, there are only trade-offs; and you try to get the best trade-off you can get, that's all you can hope for."

Thomas Sowell

This chapter will extend the Surrogate Optimiser (SO) methodology outlined in Chapter 4 and implemented for self-interested agents in Chapter 5 by extending the objective function for agent retrofit decisions to include a measure of environmentalism. This is in order to acknowledge and account for the observed tendency of individuals to be willing to pay to reduce their own greenhouse gas emissions, the evidence for which was discussed in Section 3.3 of the related works. This extension of the SO method will improve the descriptiveness of the model to capture observed aspects of reality and allow the analysis of policies that attempt to encourage environmentalism within a population.

The structure of the chapter will be as follows. Section 6.1 will explain how emissions have been implemented into the workflow, including embodied emissions, lifecycle costs, and models of national grid decarbonisation which make up the Lifecycle Carbon Savings (LCS) model. Section 6.2 will then explain how the Genetic Algorithm implemented in Section 5.3 was adapted to generate the solutions for agents of heterogeneous preferences, followed by the results of this process on a sample of households. Section 6.3 will then explain the process of training the multi-objective SO using that sample and evaluate its performance against an unseen test set. Section 6.4 will then expand the simple agent-based model introduced in Section 5.5 to take advantage of the transparent measure of household environmentalism integrated into the SO models to evaluate the performance of a simple government campaign to boost households' environmentalism. Section 6.5 will conclude the chapter with some points of discussion relating to the ideas raised.

It will be noted that there are some similarities between the chapter outline above and the structure of 5. This is unsurprising, as the expansion of the SO into multiple objectives to account for environmentally conscious agents will follow the same methodological beats as the single-objective SO did. It should be noted, however, that the energy simulation and subsequent SEPM that were implemented and trained in the previous chapter are used in this method unchanged. This is because the underlying energy performance of buildings is unchanged by the economic and social forces that influence the retrofit decision makers. Any reader wishing to understand the method used for calculating energy use this chapter should therefore refer back to Section 5.2.

6.1 Implementing Emissions

The two core components that make up the emissions profile of a dwelling are the embodied emissions and the operational emissions. The embodied emissions represent the emissions required to implement the retrofit through the manufacture, transportation and disposal of the materials and byproducts of the retrofit installation, while the operational emissions refer to the emissions produced from the dwelling as it is used over time [199]. In this instance, operational emissions are produced from the fuel sources required to meet the heating demand of dwellings. In the case of gas, this occurs at the point of use, while electrical energy is generated in a variety of different forms which are combined to form the national grid.

Given the phenomenon of interest is the greenhouse effect, understood as the impact of certain environmental emissions on the global temperature profile, we have used an equivalency measure for emissions throughout the research. Emissions have been converted into tonnes of carbon dioxide equivalent (tCO2e). This converts non-carbon gasses with a greenhouse effect to give a single unit measurement of the effects and it is a common and accepted universal emissions measure used by researchers and governments to evaluate the environmental impact of greenhouse gas emissions [200, 201, 202]. It should be noted, therefore, that the environmental impact considered in this chapter is only limited to the effect of emissions as it pertains to carbon-equivalent units and does not factor in environmental concerns such as clean water, particulate emissions, or biodiversity, which should be considered out of the scope of this research. This is considered a reasonable limitation, however, as the retrofit measures under consideration are not generally associated with significant causes of localised pollution of these kinds.

$$LCS = \sum_{t=1}^{t=p} e_t - E_0 \tag{6.1}$$

Equation 6.1: Lifecycle Carbon Saving (LCS) calculation.

Lifecycle Carbon Saving (LCS) is represented by Equation 6.1. The value of emissions-saving in year t, e_t, will depend mostly upon the energy saving properties of the retrofit and was discussed in Section 6.1.2. It will also be impacted by fuel type, as we model a decarbonising grid resulting in decaying emissions contributions from electric heating methods over time. The lifetime embodied carbon is represented by E_0 and considers the cost of manufacturing, transporting, and disposing of building materials used in a given WHRS. This can be considered a non-discounted emissions equivalent of the NPV calculated and is expressed in tCO2e.

6.1.1 Emissions Cost Model

Embodied emissions represent the emissions required to install a given retrofit through the materials and services required to install it [203]. The inclusion of embodied carbon was deemed important as low carbon buildings tend to embody a greater share of whole lifecycle carbon in their construction, despite lower emissions overall [204]. This is particularly important when considering retrofits to existing stock that is not at the end of its life, as this carbon cost can be mitigated entirely through inaction. Embodied emissions values were taken from the Inventory of Carbon and Energy (ICE) database where available, with interpolation used to obtain missing material thicknesses [177].

The approach to measuring the emissions impact in this research has been to, wherever possible, use lifecycle emissions values for the given installations. This mirrors the approach used in the financial modelling with the exception that no discounting rate has been placed on future emissions. When referring to lifecycle emissions for a given installation, this refers to the carbon not just required to manufacture the material but also to transport, dispose and install it.

As well as improving the accuracy of any carbon savings values, the inclusion of embodied carbon is conceptually important for the optimisation stage. Without embodied carbon calculation, the emissions reduction objective would become unbounded and agents wishing to minimise emissions would always install the maximum insulatory measures regardless of their diminishing energy mitigation returns. The inclusion of embodied emissions creates a tipping point at which the marginal emissions reduction from energy insulation is exceeded by the embodied carbon from additional thickness. This bounds the optimisation problem and therefore makes it solvable.

6.1.2 Emissions Mitigation Model

$$S = \sum_{t=0}^{N} f_{et}c_{et} + f_{gt}c_g \tag{6.2}$$

Equation 6.2: Lifecycle carbon reduction over N years. The net emissions per unit of electricity in year t given as c_{et} and the net emissions of natural gas are assumed fixed at c_g. The change in demand (for year t) of electricity and gas as a result of the retrofit as f_{et} and f_{gt}.

Emissions mitigation is the gross reduction in greenhouse gas (GHG) emissions as a result of a given WHRS. In this instance, all reductions in emissions will be due to a fall in operational emissions as the net embodied carbon of a preexisting building can be considered zero at the point of decision. Within the modelling scope, these operational emissions can be reduced in two ways: a reduction in energy use, or a change in energy fuel source to a lower emission producing source. This is represented in Equation 6.2 which represents the annualised emissions mitigation of a given WHRS over a period of N years. This calculation is similar to the cost-savings calculation used in the financial model, with the primary difference being the lack of discounting of future emissions compared with present emissions. The rate of GHG emissions from natural gas fuel sources is set at a fixed rate, while the electrical grid is assumed to vary based on a grid decarbonisation model to capture the reduction in expected emissions over the period of interest as the national electrical grid decarbonises.

6.1.2.a Grid Decarbonisation Rate

$$r = 1 - \frac{c_t}{c_0}^{\frac{1}{t}} \tag{6.3}$$

Equation 6.3: Decarbonisation rate required to achieve CO2 levels of c_t after t years. An inversion of the compound annual growth-rate formula.

It is the stated policy of the UK government to decarbonise the UK power grid by the year 2050 [205]. Therefore when considering the carbon emissions of power from the grid, a decay factor was considered to ensure the full benefits of the decarbonised grid are factored in. This also means that optimisations under these scenarios have an a priori assumption that the national grid will meet its policy statements. To calculate the required decarbonisation rate, the formula in Equation 6.3 is used. Given 2018 carbon levels (c_0) of 0.309 kg/kWh [6], falling below 0.001 kg/kWh (c_t) by 2050 (t = 32) the required annual reduction must be 16.4%[1]. The decarbonisation rate is highly relevant when considering extremely carbon averse households, as when observing a highly decarbonising grid they would be more likely to consider using more expensive electric heating even when gas central heating is available.

6.2 Multi-Objective Optimisation using a Genetic Algorithm

6.2.1 Objective Function

In order to incorporate the environmentally conscious agents, Lifecycle Carbon Saving (LCS) was added to the agent utility function. There were two primary objective values considered during the scoring stage of the GA. LCS represents the CO_2e saved by energy demand mitigation, less the embodied carbon cost of the retrofit. Both objectives consider lifetime values, including transportation, labour, and, where relevant, disposal

[1] Using a reduction rate formula, a value of 0 would not be achieved in a finite time period, so 0.001kg was used as an approximation.

of the materials used, as seemingly beneficial retrofits can become unfeasible if other lifecycle costs are considered. Where applicable lifecycle maintenance costs were rolled into upfront costs with appropriate discounting applied. While the assessment of the benefits of a retrofit assume it will be operated for its full lifespan, there is no conceptual reason that a second retrofit could not be applied during this period given the initial installation is a sunk cost.

The household utility for a retrofit solution is shown in Equation 6.4, with the household's carbon valuation expressed as v. Given that v is expressed in £/tCO2e, the utility can be expressed in financial terms, allowing a direct trade-off between the terms. This carbon valuation, v, represents the Willingness To Pay (WTP) for carbon mitigation and is discussed in terms of the Household's Carbon Valuation (HCV). The driver for this value can be seen as the personal or social components that motivate households to be carbon-conscious and could be driven by a combination of altruism, social pressures, or self-preservation. A discussion of the conceptualisation of HCV can be found in Section 6.5.1

$$U = \mu NPV + vLCS \tag{6.4}$$

Equation 6.4: Utility function used in optimisation, with v representing Household Carbon Valuation (HCV). μ is set to 0 for a LCS-maximising objective but is otherwise 1.

$$V_{u0} = \frac{-NVP}{LCS} \tag{6.5}$$

Equation 6.5: Minimum indifference carbon value.

Equation 6.5 shows the minimum indifference carbon value. This is the minimum carbon value required for a household to be indifferent between inaction (their existing building state) and a given retrofit solution. If a household's individual carbon value exceeds this threshold, and the WHRS has a positive LCS, a utility-maximising household would benefit from adopting the solution. This value has some useful properties, being both unique for Pareto solutions and monotonically increasing on a concave Pareto front. It should be noted that the minimum indifference carbon value differs from the HCV considered at the point of optimisation. To see this, consider that in positive NPV solutions which also have positive LCS (a subset that makes up the majority of the WHRSs observed in Section 5.3.3) the minimum indifference carbon valuation is negative despite being acceptable to an NPV-maximising household with a HCV of £0/tCO2e.

6.2.2 Pareto Front Generation

Pareto fronts represent the subset of all possible solutions for which there exists no other solution that performs better on all objectives. A background discussion of this concept was presented in Section 2.3.2. In the case of the two objective optimisation problems under consideration, this means that if a WHRS is part of the Pareto front then no other solution exists that has both a greater NPV and LCS. In order to train the SO, a representative set of HCV-linked Pareto fronts was required. Pareto solutions were generated by altering the HCV provided to the objective function. HCV bounds were generated by using a NPV-and-LCS-only objective to find the extreme points of the front. A grid search was then performed between these bounds and a local search was performed around the identified solutions. An example of the type of front generated by this method is shown in Figure 6.1. Using this front generation technique, in contrast to a multi-objective GA method, provides a one-to-one relationship between the HCV and optimised solutions for preference-aware SO training.

6.2.3 Results

Using random sampling, a total sample of 2,635 buildings were selected to be optimised, from which a total of 45,175 WHRSs were discovered, resulting in an average Pareto front size of 17.1 solutions. The sample size was dictated by a given computational budget[2]. This section will investigate the nature and distribution of WHRSs found, both to better understand the preferences of environmentally conscious agents from this sample and to investigate the training set to be carried forward into SO training.

In order to analyse the WHRSs generated they will be broken down into three scenarios. The Max NPV scenario mirrors the problem discussed in Chapter 5, with HCV set to £0/tCO2e to capture the decisions of purely self-interested agents. The inverse scenario of this, Max LCS, represents solutions for which μ in objective function 6.4 is set to 0 and agents attempt to maximise LCS regardless of the financial implications. And solutions between these two scenarios will be considered part of the Mixed Criteria scenario. These solutions have a μ of 1 and a positive HCV (v) and represent all households who place a positive value on carbon mitigation while remaining interested in the financial returns those WHRSs yield.

[2] 400 cpu hours were assigned to this instance of WHRSs discovery.

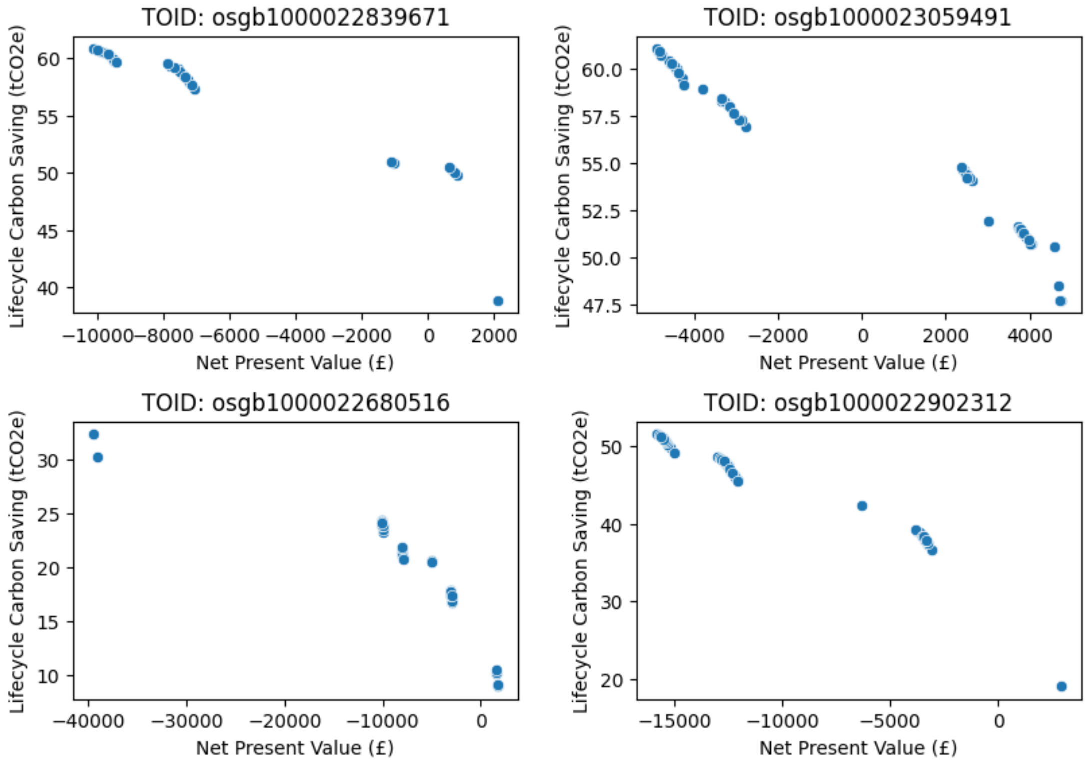

Figure 6.1: Sample of Pareto fronts generated with the environmentally conscious objective function and the Genetic Algorithm.

An initial investigation into some representative fronts will give an idea of the nature of the WHRS discovered. Figure 6.2 shows a typical Pareto front from a sample dwelling. Solutions can be seen clustered around general intervention type combinations, with local non-dominated alternate solutions resulting from gradual variation in insulation thickness. In this instance, the positive NPV solutions were only all single-measure IWI installations, although it was observed in Section 5.3.3 that a variety of WHRSs could be discovered for different dwellings even when restricted to profitable combinations. The next cluster of discovered solutions was the installation of triple glazing alongside the IWI, followed by the addition of supplementary roof insulation, then by the replacement of EWI with IWI at higher carbon valuations. There is variation in the nature of the fronts as the characteristics of each building differ significantly. In the dwelling with the TOID: osgb1000022680516 for example, which can be seen at the bottom left of Figure 6.1, there exists a set of solutions with significantly negative NPV. These are WHRSs in which the heating system is converted from gas to electric in order to capitalise on the decarbonising electrical grid discussed in Section 6.1.2. This is the inverse of the transition observed in many NPV-maximising solutions in which electrically heated dwellings converted their energy generation to gas due to the significant fuel cost savings. It should be noted that this transition is very costly, with a minimum indifference carbon value of £1300/tCO2e in this instance. This is well outside of the generally discovered WTP for emissions mitigation and these solutions were generally only found in the Max LCS scenarios.

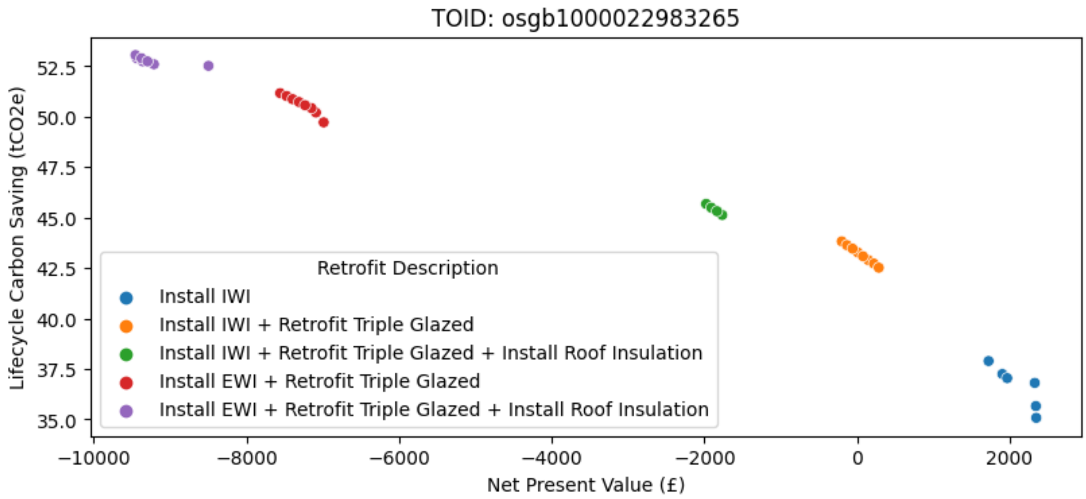

Figure 6.2: Sample Pareto front coloured by retrofit installation.

A breakdown of the entire sample of WHRSs demonstrates that the trends observed in the representative front can be seen across the rest of the samples. The solution trade-offs between the two objectives are plotted in Figure 6.4 which is coloured in accordance with the retrofit measures installed. As noted above, the most extreme solutions are made up of fuel source retrofits, with electric-to-gas resulting in a high NPV but low or negative LCS while gas-to-electric causes the opposite trend. This demonstrates that the grid decarbonisation is significant enough for heavily carbon-averse households to eschew the cost advantages of gas fuel in order to reduce emissions in this case. The same wall insulation trend observed for the representative front was also seen across the sample, with IWI resulting in a higher NPV due to its lower cost but EWI dominating the LCS-focused solutions. This is perhaps unsurprising given the higher cost but additional insulating effects of EWI. The glazing follows similar trends, with double glazing popular in WHRSs focused on NPV while triple glazing is more common in solutions with a lower or negative NPV but a higher LCS.

The distribution of optimal thickness of insulation types varies between scenarios, as can be seen in Figure 6.4. The Max NPV scenario distribution mirrors that discussed in the previous chapter, with a fairly even distribution of wall insulation and clusters of roof insulation towards the extreme ends. The Max LCS scenario finds maximum roof insulation thickness to be universally optimal, indicating that the embodied carbon of the insulation is more than offset by the resultant energy savings. In contrast, approximately 30% of wall insulation installations in the Max LCS scenario are sub-maximal, indicating more significantly diminishing returns on the embodied carbon as thickness increases. As would be expected, the mixed criteria scenario is an intermediate distribution between the Max NPV and Max LCS scenarios, with a more even insulation thickness distribution than the LCS-maximising solutions but a steeper and more right skewed distribution of both wall and roof thicknesses than the NPV-maximising solutions. This effect is more significant for roof than wall insulation, adding evidence that more emissions-focused agents will focus more on maximal roof installations while seeking sub-maximal wall insulation.

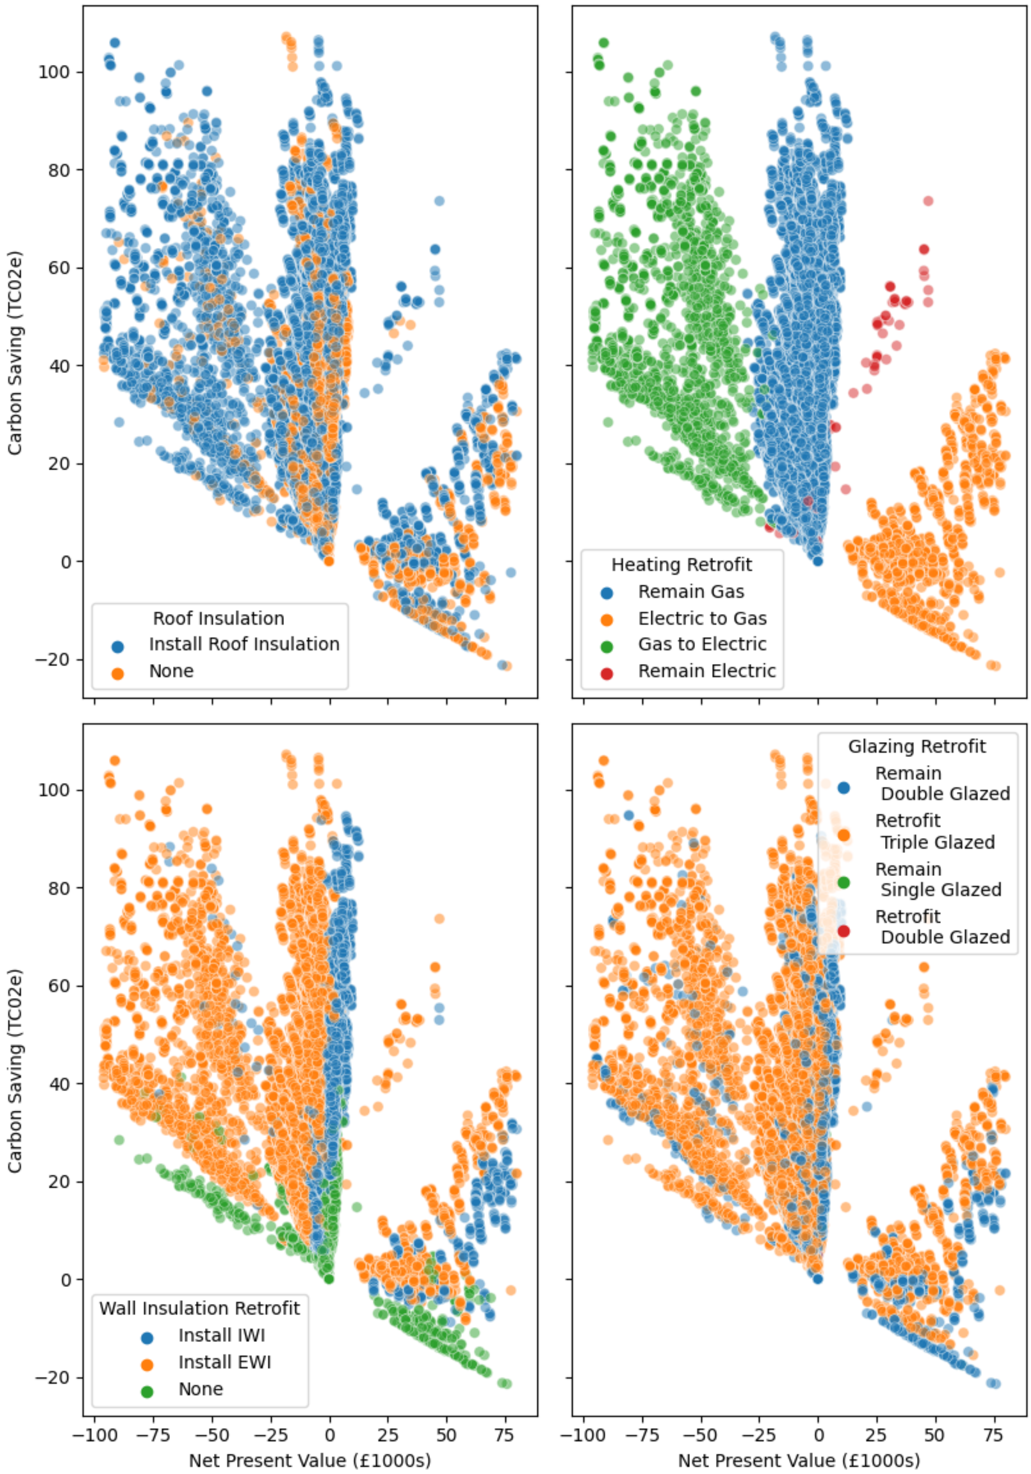

Figure 6.3: Entire sample of 45,175 WHRSs trade-offs coloured by retrofit measure.

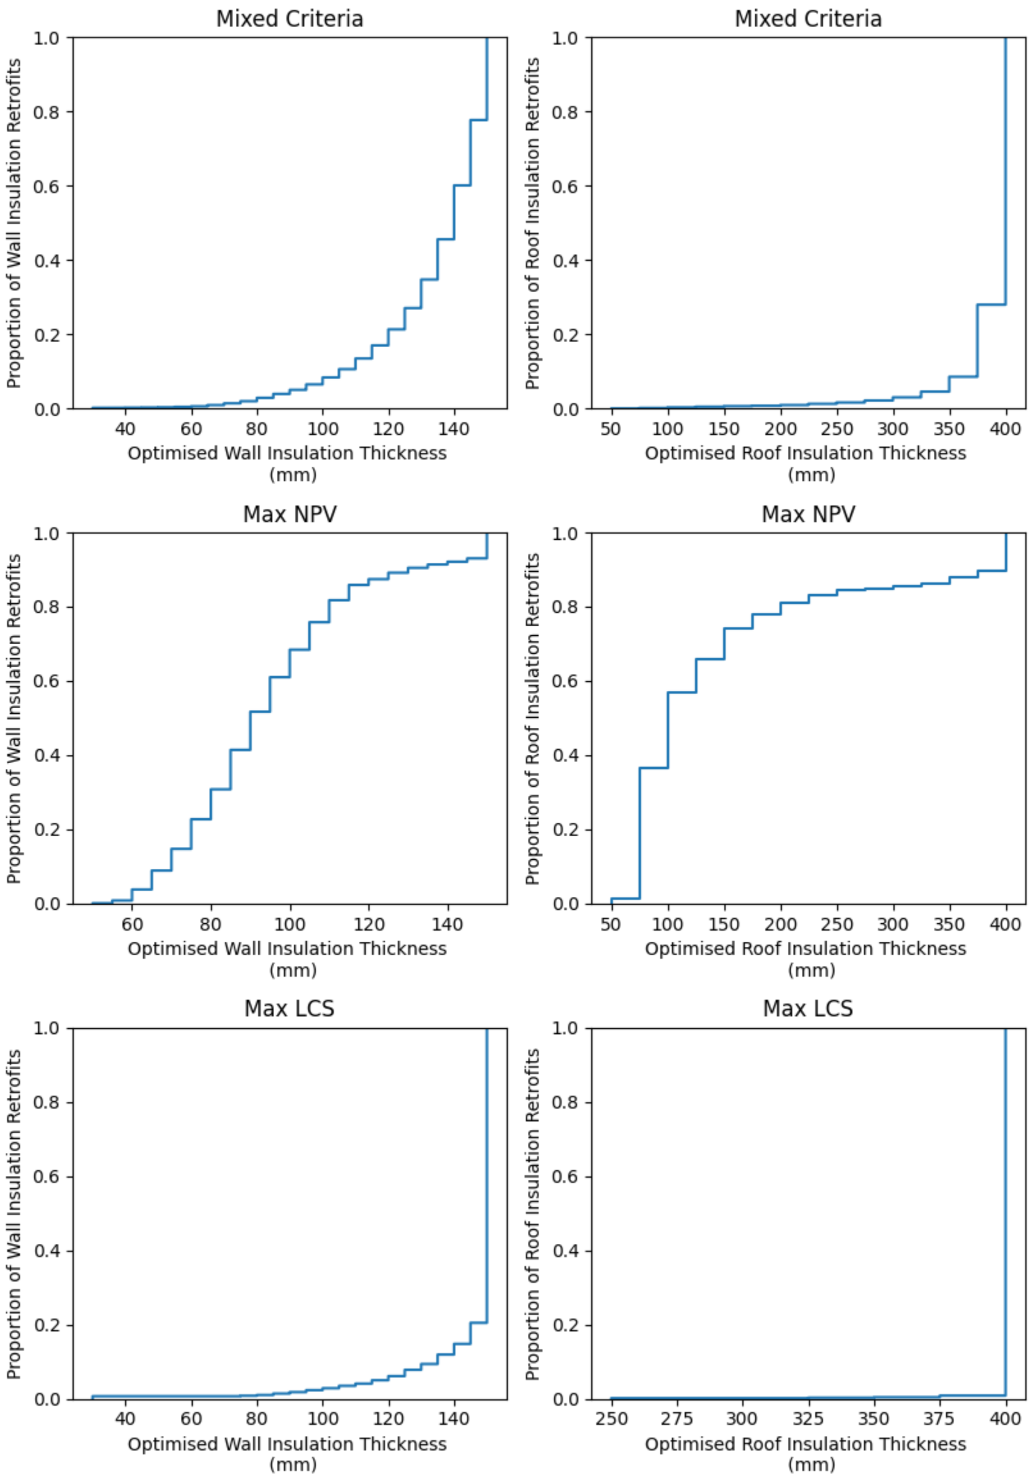

Figure 6.4: Distribution of insulation thicknesses broken down by scenario.

6.3 Multi-Objective Surrogate Optimisation

At this stage, the Pareto fronts for 2,635 dwellings have been discovered for a total of 45,175 WHRSs; These solutions are dependent on both the initial state of the building and the carbon preferences of the household that are used to generate them. The same fundamental method described in Section 4.4 and implemented in Section 5.4 was used to train the set of predictive ANNs that comprise the SO, with the addition of the HCV as an input feature to account for the environmental preferences of households.

6.3.1 Model Training

The hyperparameter tuning procedure was generally the same as described in Section 5.4.1. However, some of the elements were adjusted based on initial observation of training outcomes. The Hyperband algorithm used for single-objective SO parameter tuning was found to be too aggressive in selecting models which performed well initially but were insufficiently complex to capture the impact of HCV on model predictions and would therefore perform poorly. To compensate for this, a Bayesian optimisation algorithm was implemented for parameter tuning instead. This algorithm applies a Gaussian process, calculating a prior distribution, performing model training and updating the posterior based on the results [206]. This optimises hyperparameters efficiently while avoiding some of the early over-fitting problems experienced using Hyperband. The range of hyperparameters over which the models searched can be seen in Table 6.1.

Setting	Regression Models	Classification Models
Loss Function	L2 (Least Square Errors)	Categorical Crossentropy
Last Layer Activation	Linear	Softmax
Pre-proccessing	Normalisation, Logarithmic, Standardisation	
Number of Layers	1 - 7 (Steps of 1)	
Neurons Per Layer	8-256 (Steps of 8)	
Activation Function	ReLU	
Learning Rate	0.01, 0.001, 0.0001	
Training Function	Adam (alpha=0.001, beta1=0.9, beta2=0.999, epsilon=10E-8)	

Table 6.1: Description of SO hyperparameters used for tuning.

Generally, the models were trained using 80% of the whole data set, with 10% retained for testing, and 10% for validation. However, in order to ensure the input data was balanced and matched the application of the trained model, the IWI and EWI thickness models were trained only on the data set of WHRSs for which an installation of this type was made. This was to ensure the significant number of zero valued thicknesses that comprise the entire set did not bias the model during training. The classification models were trained on the entire set, ensuring that correct classification would result in subsequently correct thicknesses.

6.3.2 Feature Analysis

In order to ensure that the carbon value used to derive the choice of retrofit prediction, an important feature if the value is to be meaningfully applied to agents in predicting decisions, a form of feature analysis is required. SHAP (SHapley Additive exPlanations) is a form of additive feature analysis first introduced by Lundberg and Lee in 2017 [207]. This game theory approach considers a weighted power set of feature combinations to calculate the impact of each feature on the predictions made by a model. Larger SHAP values associated with a feature indicate that the inclusion of that feature has a larger impact on the predictions of the model. A good primer on the SHAP method can be found here [208]. The SHAP python implementation was used to obtain mean SHAP values [209].

The SHAP values are computationally expensive to gather and as such a sample of 500 dwellings were used to calculate the significance of each feature on each model. The results showing the most contributing features to each model are shown in Figure 6.5. The mean absolute SHAP values have been used to gather the magnitude of impact on the model, although the values are not plotted as the relative significance of features is of interest rather than an interpretation of the absolute SHAP values.

The level of influence that each feature brought upon the model varied between the measures over which the optimisation occurred. There are some expected trends in the data, with HCV being a more prominent feature in the measures that varied most over the Pareto fronts seen in Section 6.2.3. IWI material selection, for example, saw variance across the Pareto front from PIR for low HCV households, then to EPS, and finally to None as EWI installations were priorities by highly environmentally conscious households. This is reflected in the significant influence of HCV on the classification of those particular materials, as seen in Figure 6.5. Across all the models, HCV is a significant feature, ranging from the single most contributing factor to the 5th, indicating the models are highly influenced by the preferences of individual households.

Examining the most significant features of the individual models increases the interpretability of the ANN models and allows for some degree of validation. It is reassuring for example that the total area, wall area, and initial energy use are significant features in determining an appropriate IWI or EWI material, as these are the most significant contributing factors to both the cost and energy models that make up the underlying optimisation problem. Interestingly, the wall type is a more significant feature in IWI selection than in EWI selection. When considering the sample of near-optimal retrofits used for training this makes intuitive sense, as the IWI material sample set draws from the lower HCVs and, therefore, higher NPV seeking preferences. In this case, the existing wall type would contribute a lot to the profitability of the installation, as insulated cavity walls provide significantly less heat loss than solid walls, making the additional IWI less effective. This explains why the presence of solid walls negatively impacts the model's propensity to predict IWI insulation. This factor is less significant when modelling EWI, as the training set of WHRSs which included EWI is from the more environmentally conscious side of the Pareto front and as such, the minor difference in marginal benefit that wall type yields is less significant than the presence of other thermal properties of the house such as the roof and glazing properties.

Examining the feature significance of the glazing prediction model indicates that the most significant properties are, as may be expected, the glazing ratio of the building as well as the initial state of the building's glazing. Interestingly the side glazing ratio was brought forward as a significant feature during training despite other properties such as total glazed area or wall area. This may be because using the side glazing ratio was more efficient for approximation during training than the more general values and provided more consistency across different sizes and topologies of dwellings.

The wall insulation thickness models, both IWI and EWI, were heavily dependent on initial energy use as a key predictive feature. The HCV was the 4th and 5th most significant feature respectively, indicating the models are responsive to the preferences of the agents at hand. It could be asked whether the relative insignificance of this feature is cause for concern, compared with some of the other models. However, the training set of wall thicknesses only included those WHRSs for which insulation of that type was installed. Given that the IWI and EWI classification models were significantly HCV dependent, it is intuitive that the feature is of less relative importance as the training set is essentially pre-selected based on a very HCV dependent attribute (the chosen insulation type). This hypothesis is consistent with the features of significance found in the roof insulation thickness model; outputs are significantly more sensitive to the HCV feature as the presence of some roof insulation is significantly less sensitive to HCV. This was observed in Section 5.3.3 where even NPV-focused solutions (where HCV is set to 0) would prioritise some level of roof insulation. This means the roof insulation thickness model training data was not pre-screened based on HCV sensitive criteria and subsequently this feature is of much higher importance.

6.3.3 Surrogate Optimiser Results

A test set of 4,637 WHRSs was used to measure the performance of the trained SO models. The results of the classification models can be seen in the confusion matrices of Figure 6.6, a histogram of regression residuals of the thickness models can be seen in Figure 6.8, and a cumulative density function of the predicted thicknesses against the test data set distribution is shown in Figure 6.7. This section will discuss the performance of these models individually.

The model classification of IWI material installations performed well, with a weighted F1 and accuracy score of 0.90. The worst performing classification was of XPS installation, with a recall of 0.71 and a precision of 0.75. The similarity between the precision and recall values indicates a similar level of false positives and false negatives, which can be seen in the confusion matrix in Figure 6.6. Classification errors were also balanced between the classes, with a proportional number of misclassifications between XPS and PIR compared with XPS and None. The EWI classification model also performed well, with an accuracy of 0.89 and a weighted F1 score of 0.87, indicating generally good performance. The PIR classification in the model performed poorly due to the significant imbalance of the classes. However, the absolute number of errors was quite small. Additionally, the misclassified PIR values are generally classified as XPS, which will be a significantly smaller source of final score error compared with them being assigned to None. It is therefore clear that while the model is unable to distinguish sufficiently between the two alternative materials, it is sufficient to perform a binary classification to determine where some form of insulation is required.

The glazing classification performance is good, with an accuracy and weighted F1 score of 0.91 and 0.92 respectively. Most observations fall into either double or triple glazing and this is represented in the errors, which are fairly evenly distributed between false positives and false negatives in double glazing predictions as shown in Figure 6.6. Given the larger support for triple glazing, this resulted in poorer metrics for double glazing of 0.88 and 0.85 for precision and recall respectively.

The fuel source classification model performed well, correctly classifying 285 of the 295 instances of gas to electric conversion in the test set with only a single instance of a gas fuel source dwelling being incorrectly classified as electric. The remainder of dwellings were correctly classified as optimally utilising gas for heating, generally preferred due to the significantly lower economic cost. This was one of the best performing classification models, likely due to the significant HCV required under this scenario for the fuel type switch to increase utility due to the significant cost differences making only the most LCS-focused solutions candidates for the model. The quality of these classifications was represented in an accuracy of 0.98 and a weighted F1 score of 0.99.

The classification scores are similar to those achieved in the single-objective SO problem presented in Section 5.4.3. The regression problems also performed well, although the MAE was slightly higher than in the NPV-maximising case due to the increased problem complexity. The IWI thickness model resulted in a MAE of 8.4mm, slightly larger than the discretised thickness class sizes which were spaced 5mm apart. Nonetheless, the modal error band is still 0mm, as shown in Figure 6.8, which shows the distribution of regression residuals. The final distribution of IWI predictions maps closely to those in the test data set, making the model suitable for the desired upscaling tasks. The distributions of the regression model predictions are plotted against the test set in Figure 6.7. The EWI thickness model performs slightly better with a MAE of 7.8mm which again represents slightly more than the size of a single class range when coerced back into discrete class values. The roof thickness model performs worse in absolute terms with a model MAE of 10.3mm, however, the roof thicknesses occur over a wider range and with wider class bins of 15mm and so the mean error falls below the width of a single class. The IWI, EWI, and roof insulation thickness models obtained r^2 scores of 0.76, 0.83, and 0.88 respectively, indicating that in all cases a significant proportion of the output variance is explained by the model.

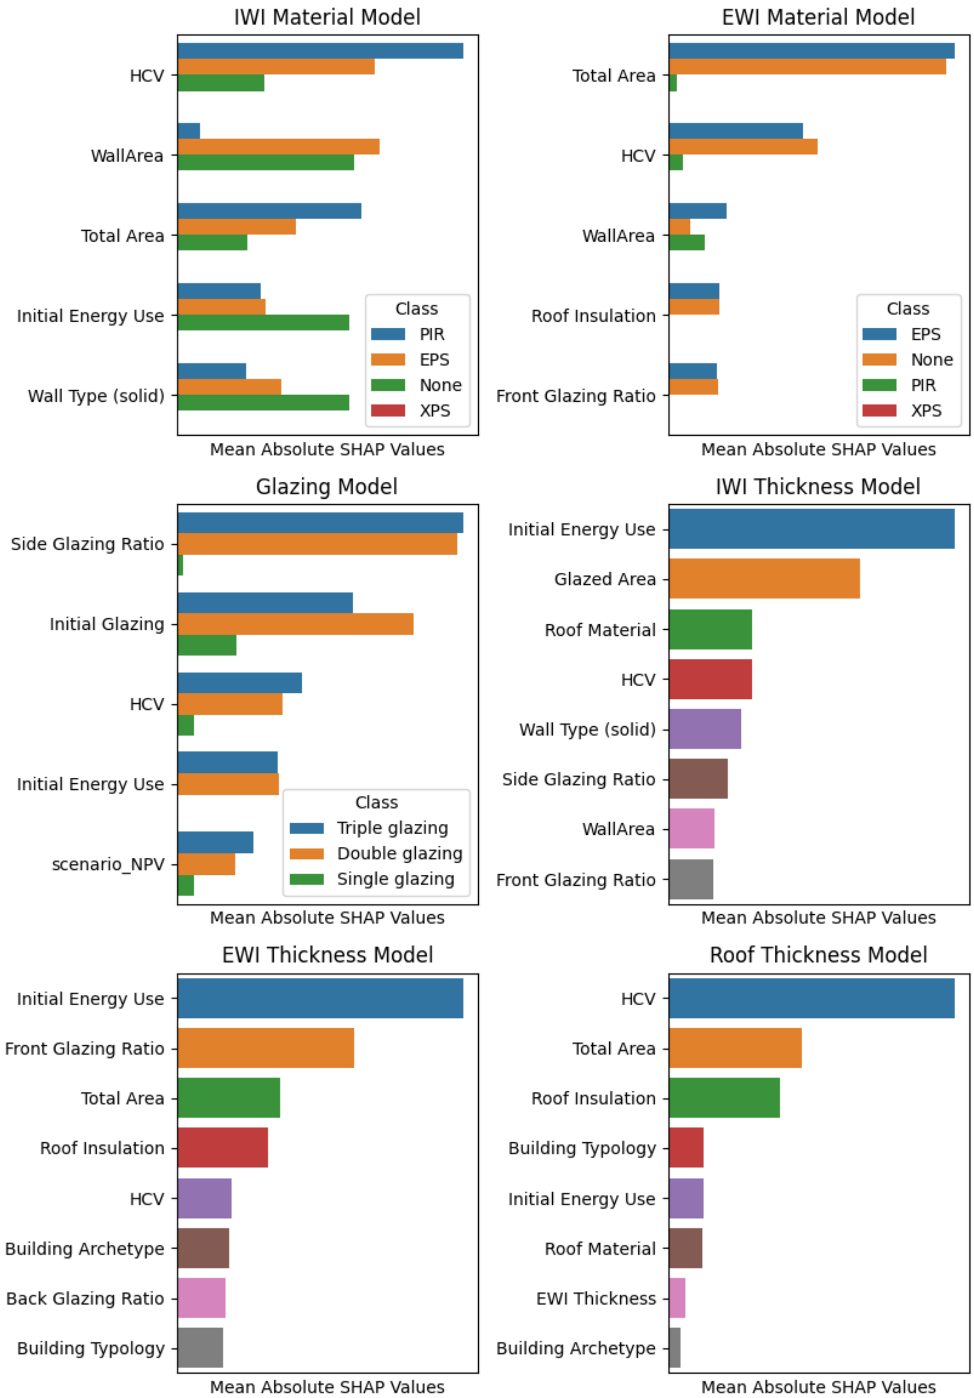

Figure 6.5: SHAP values indicating the significance of different features to model predictions.

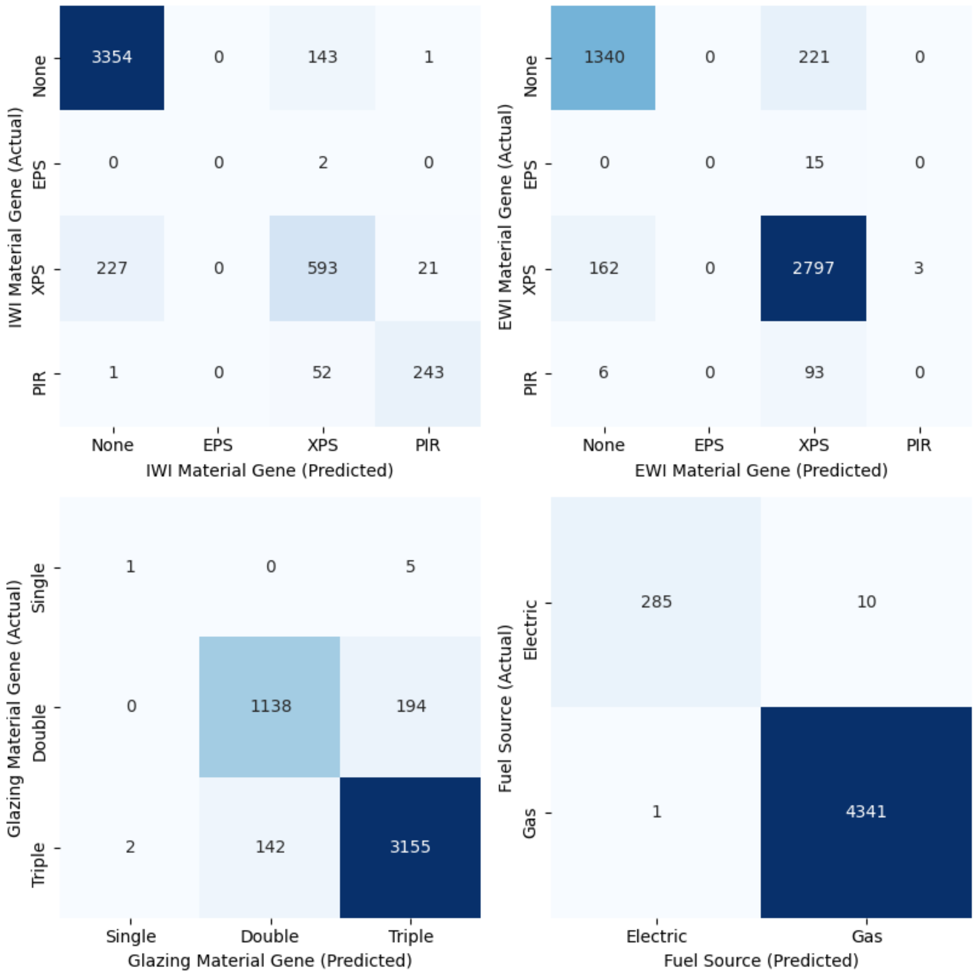

Figure 6.6: Confusion matrix of classification predictions.

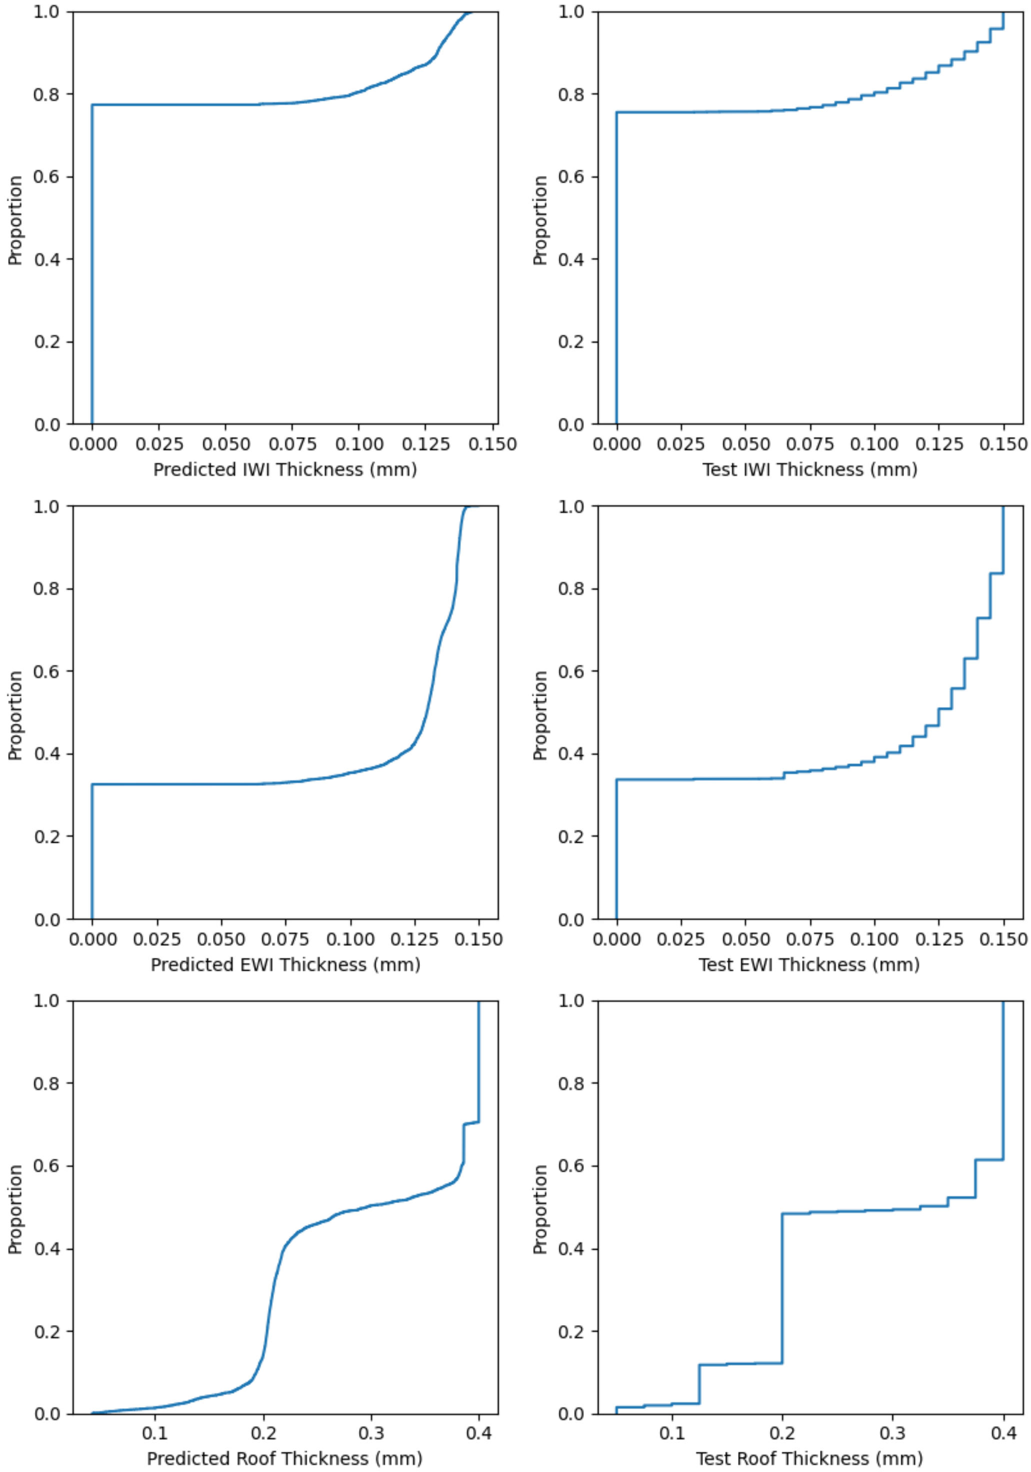

Figure 6.7: Cumulative density functions of actual and predicted insulation thicknesses.

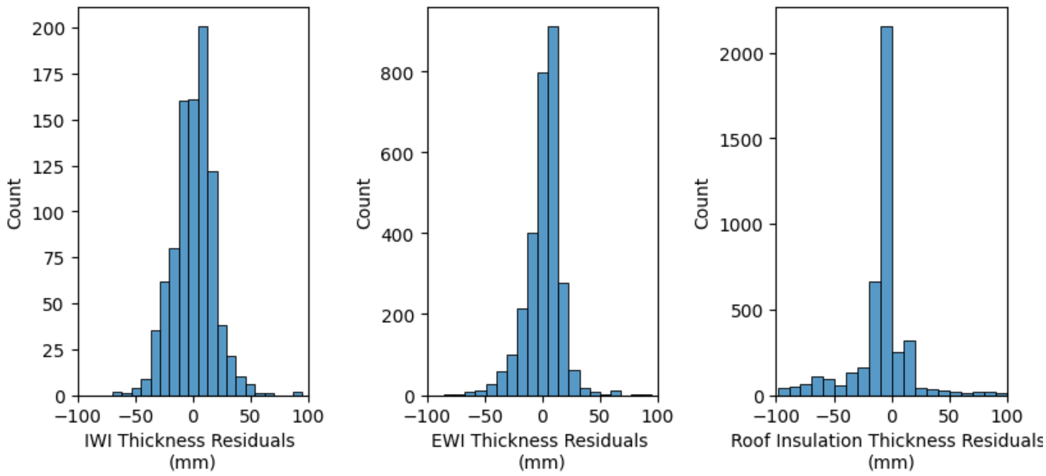

Figure 6.8: Histograms of regression residuals for thickness predictions.

6.4 Agent-Based Model of Retrofit Adoption using Multi-Objective Surrogate Optimisation

In order to demonstrate the benefits of adding Household Carbon Valuation (HCV) as a transparent measure of environmentalism into the SO, the ABM used in Section 5.5 has been extended. The illustrative ABM allows the modeller to assess the impact of a local government campaigning for climate awareness on the quantity and quality of retrofits within the city. This section will first layout changes to the model design which allow transparent integration of HCV into the model. In the results section, the results of several experiments will be presented to both observe the impact of incorporating environmental agents and evaluate the effect of variation in the model parameters.

6.4.1 Model Design

In the model, agents are endowed with a HCV of $v_i \sim \mathcal{U}(0, 50)$ drawn uniformly between 0 and 50. The government agent in the model runs a campaign to boost the environmental awareness of each household that is influenced. The rate of influence on the households will be referred to as the penetration rate p, and refers to the probability that a given household will be influenced within a given year. If influenced by an environmental campaign, the HCV of the household is increased by $n \sim \mathcal{U}(0, 15)$. The pre and post campaign effect on emissions valuations represent approximations of the mean and variance of WTP for mitigation found in carbon valuation studies [145, 146]. However, they should be considered parameters that may need localising or adjusting based on the modeller's expectation of the target system.

As with the original model presented in Section 5.5, agents evaluate their retrofit options at a probability e per year. The household's near-optimal retrofit solutions are generated using the trained SO and scored as described in Figure 6.10. Solutions bearing positive utility are then installed within the model. The modelling scenarios give the local government two approaches to influence retrofits: increasing the intensity of environmental campaigning, and direct outreach to encourage the consideration of retrofits among households (through means such as payback finance structures, working with installers, and contacting households, which are abstracted into the evaluation rate).

The state diagram representations of the agents are shown in Figure 6.9. With a domestic building stock of 95,500 dwellings, the outreach scenarios require an average of 47,750 optimisations per 25-year simulated run, representing a large computational cost that

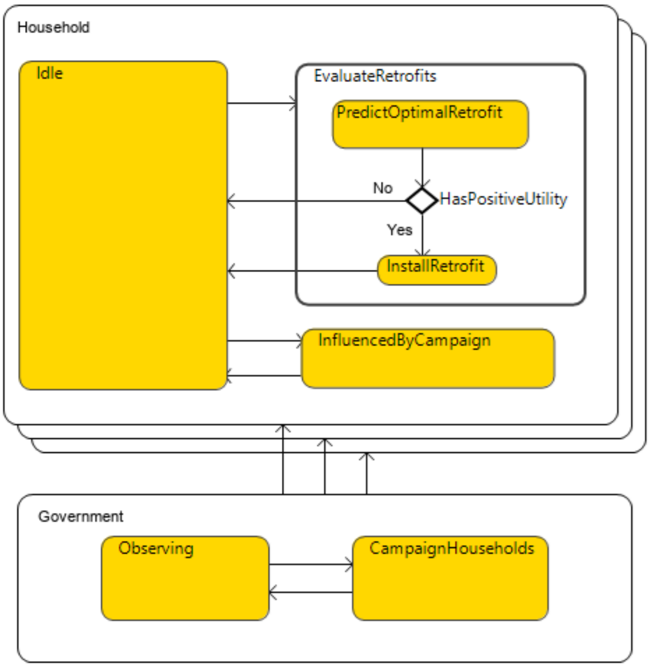

Figure 6.9: State chart representation of agents.

could become prohibitive using traditional optimisation approaches.

6.4.2 ABM Results

Scenario	Evaluation Rate (e)	Campaign Penetration Rate (p)
Baseline	0.01	0
Small Campaign	0.01	0.05
Small Campaign with Outreach	0.02	0.05
Significant Campaign	0.01	0.1
Significant Campaign with Outreach	0.02	0.1

Table 6.2: Intervention scenarios for illustrative ABM.

The scenarios modelled are shown in Table 6.2, combine the outreach principle discussed in Section 5.5 alongside the newly introduced campaign model penetration rate to create a small set of scenarios to evaluate initially.

The ABM utilised the ability of a trained SO to quickly evaluate the entire building stock of the city of Nottingham. The ABM was replicated 100 times per scenario. As shown in Figure 6.11, the campaigning to increase the HCV of households increased the

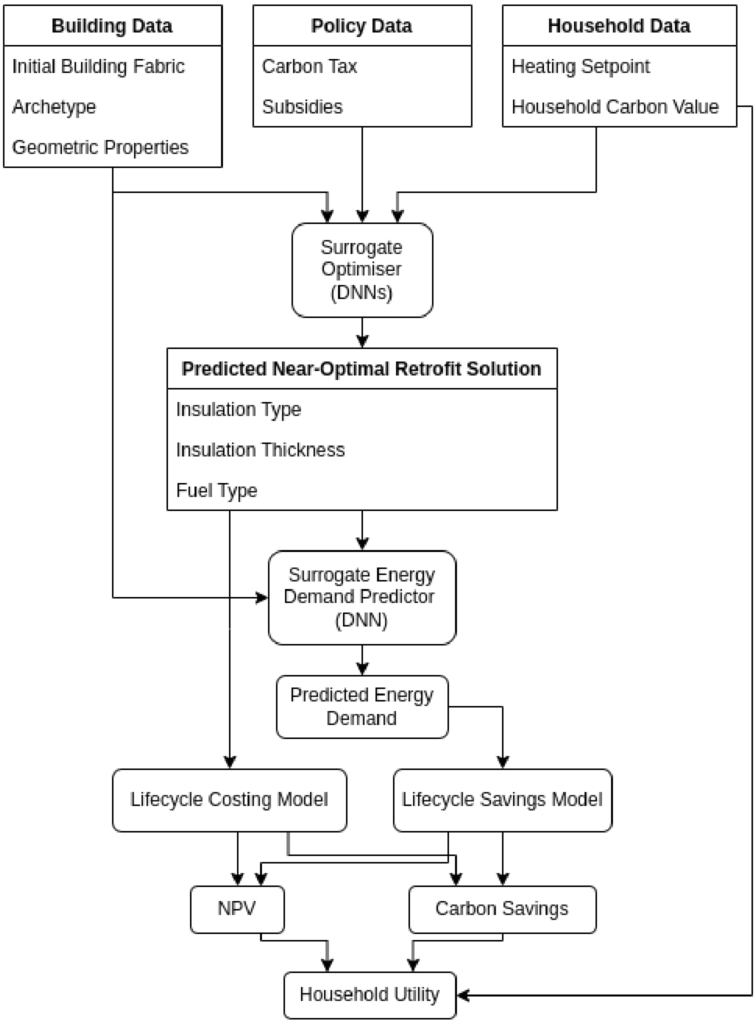

Figure 6.10: Process flow for generating and scoring retrofit solutions using the trained Surrogate Optimiser.

LCS of retrofit decisions made over time. The types of retrofits installed differed in the campaign scenarios, as shown in Figure 6.12, with increased campaigning resulting in a preference for more emissions-friendly retrofit options. This resulted in an average LCS of $25.51tC0_2e$ per retrofit in the significant campaign scenario, contrasting with $25.21tC0_2e$ in the baseline ($p < 0.01$). This came at the cost of a reduction in NPV from £7,696 to £7,429, showing a willingness of households with higher HCV to trade-off financial value for more environmental retrofits. In total, a significant campaign reduced net lifecycle emissions by $13.7ktC02_e$ without the outreach program and $16.7ktC02_e$ when outreach was included. The number of positive utility retrofits was also higher under campaign scenarios, with average increases of 2.9% and 6.3% for the small and significant campaigns respectively. This indicates the installation of marginal retrofits which would not have been viable for financial reasons alone. While the effects are small on the household level (due to the relatively moderate changes in HCV) they show that on the urban scale the increased emissions awareness can have significant impacts over a long period. In this scenario, the evaluation rate of households, influenced by the outreach campaign, had the most significant effect on the total carbon reduction. This suggests that attempts to engage households to trigger WHRSs evaluation should be a priority ahead of campaigning to increase their carbon valuation directly. However, the combined approach yielded the greatest reduction in total emissions.

6.4.2.a Penetration Rate Effect

Given the penetration rate (p) is expressed as the probability that a given household will be influenced by a campaign in a given year, the number of households affected can be expressed as a continuous value between 0 and 1. In order to evaluate the effect of the campaign's penetration, the evaluation rate was fixed to 0.02 (modelled as an outreach scenario in the previous experiment) and a parameter grid of penetration rates was performed, varying from 0.0 to 1.0 in increments of 0.1 to evaluate a full range of scenarios. Each scenario was repeated 100 times to account for the stochasticity of the model.

The results of the penetration rate experiment are plotted in a series of time series plots in Figure 6.13. As the campaign effects are cumulative over time, the scenarios are initially similar before diverging as the simulated time continues. The total number of retrofits installed increases with the penetration rate, in line with the intuition that more environmentally conscious agents have a higher propensity to install retrofits compared with self-interested agents. While this increase in WHRS installations results in higher annual investment, the total NPV of retrofits installed per year falls at higher campaign

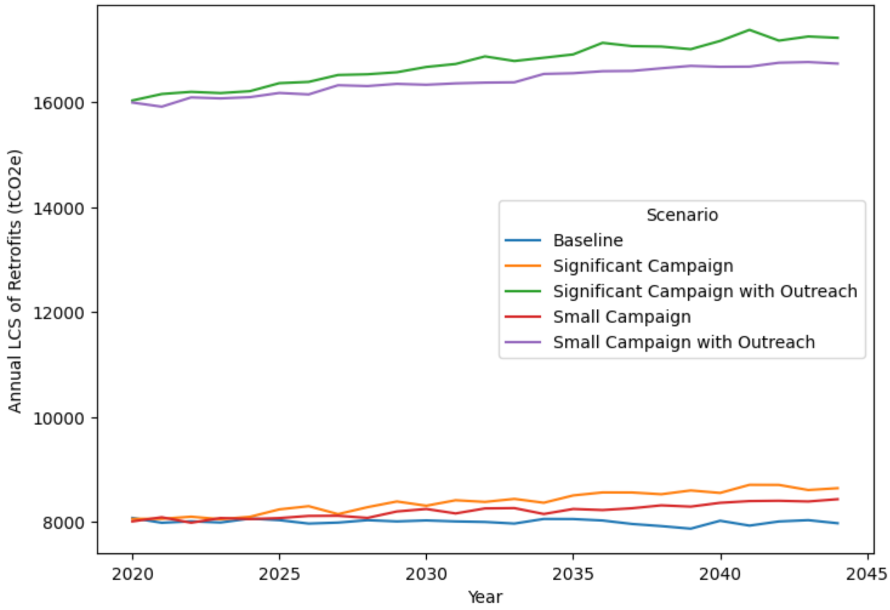

Figure 6.11: Lifecycle emissions savings of simulated retrofit decisions over time.

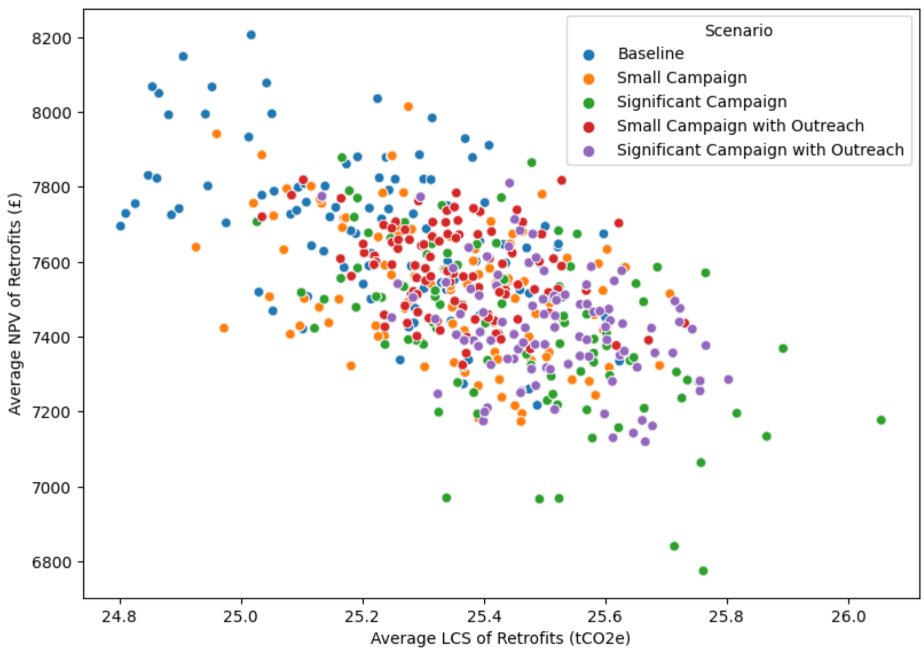

Figure 6.12: Average NPV and LCS of retrofits in the Agent-Based Model under different scenarios.

penetration rates, indicating a propensity to select retrofits with a lower or negative NPV in order to improve the environmental objectives. The total impact of a campaign with full penetration ($p = 1$) results in a cumulative GHG emissions saving of 486.7ktCO2e over the 25 year run (averaged over 100 repetitions) compared with just 447.4 ktCO2e when the campaign's penetration rate was 0. This indicates a maximum 39.5kt reduction in total LCS over the 25 year period of modelling should a campaign prove maximally effective in reaching households under an outreach scenario of 2% annual evaluation.

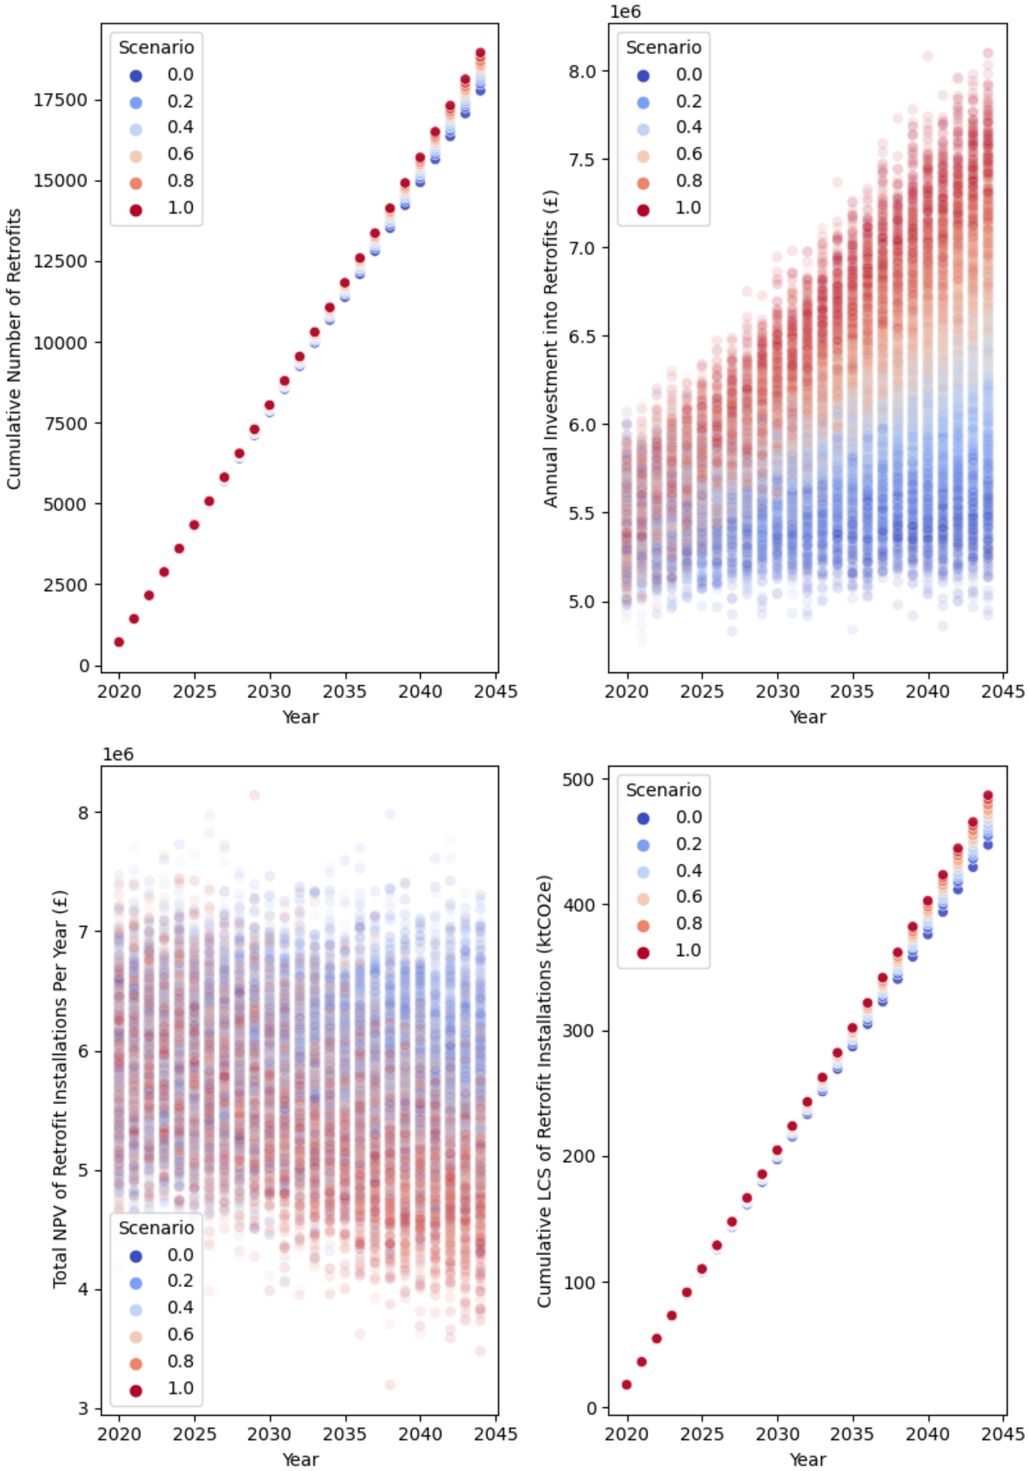

Figure 6.13: The effect over time of different campaign penetration rate scenarios on the total number of retrofits (top left), annual retrofit investment (top right), annual Net Present Value of installations (bottom left), and cumulative Lifecycle Carbon Savings (bottom right).

6.5 Discussion

6.5.1 Conceptualisation of Household Carbon Value

The primary addition of this chapter was the attribute of Household Carbon Value (HCV), which represents the value that a household is willing and able to pay per ton of CO2e that a retrofit is able to mitigate. It was observed that introducing this value made households willing to select WHRSs that are less than NPV-maximising in order to improve this secondary objective. While this notion of WTP is generally understood in a descriptive sense, it is worth discussing what the value captures about an agent's behaviour conceptually. Specifically, why might households be willing to forego financial returns in order to contribute to an improved global climate, which would be considered a public good. Section 2.4.2 of the background chapter laid out the public good game formulation in which participants have been observed to make voluntary contributions under certain conditions even when this conflicts with the Nash equilibrium predicted for the game.

The contributions could be seen as a form of pure altruism, in which participants are genuinely concerned about the impact of their carbon footprint as it related to the global climate. There are likely some issues with this conceptualisation from a rational-choice perspective, as the global climate is extremely large, and the emissions made from a single household will have a minute overall impact. Nonetheless, this is often the reason reported and often aligns with the stated intent of households. Another form of altruism to consider would be impure altruism or "warm glow" theory [52]. This would be a form of positive emotion experienced by making an altruistic decision. The good feeling of being environmental provides a positive utility similar to that which would be achieved by consumption of a good or service of the same value provided by the foregone NPV. The difference between pure and impure altruism here relates to the actual impact compared with the perceived benefit, as impure altruism is based on the agent's perception of their actions, rather than the result. In order to integrate a distinction between pure and impure sources of altruism, future work could consider omitting hidden carbon costs (such as embodied carbon) from decisions made by impure altruists, as agents are unlikely to be aware of the impacts of such elements. Other theories of altruism stem from social causes. A primary social cause of altruism would be the notion of conditional cooperation, a well established driver of public good contributions [210, 53]. This would apply to instances in which households contribute to the environment on the understanding that other households would also do so, contributing to the greater good. In a conditional cooperation scenario, the perception of others' behaviour will be a key driving force behind a household's WTP. There has been evidence presented that conditional cooperation is a dominant model in predicting environmental behaviour [211].

Ultimately, the HCV used in this chapter was deliberately designed to be agnostic to the underlying cause of public good contributions. There are likely to be multiple and varied reasons for the underlying environmentalism observed in household decisions. The prediction of this WTP on an individual household level and the root causes and models for exactly how it can change over time are outside of the scope of this research.

6.5.2 The Value of Surrogate Modelling

The stated goal of implementing a Surrogate Optimiser (SO) in this chapter was to increase the level of detail possible within a bottom-up WHRS-adoption model while keeping the model technically feasible. The introduction of environmentally conscious agents allowed households with heterogeneous environmental preferences and unique dwellings to be modelled individually at a large scale. It is worth considering whether the use of the SO procedure was able to improve the technical feasibility of the model compared with a traditional optimisation approach. A total of 45,175 WHRSs were used to train the SO, representing front values for 2,635 distinct dwellings. A single ABM run with an evaluation rate of 1% and a run period of 25 years would require 23,875 retrofit optimisations in expectation. Given the SO training time is negligible compared with generating a large set of optimisations, any modeller wishing to perform more than two ABM runs would benefit from the SO methodology when considering the raw computational cost. This is likely to be in most cases, as running just two instances of a stochastic ABM with multiple parameters of interest is unlikely. This finding holds even when considering there may be some repetition of dwellings selected across simulation runs, as even when the same dwellings are selected, the environmental preferences of the household are unlikely to be constant across runs, requiring the re-performance of the optimisation procedure. Given there is some loss of accuracy when using predictions rather than performing traditional optimisation for each dwelling, there may be cases where modellers still choose to use traditional approaches. However, this is likely to limit the scale of the modelling possible. The ABM runs presented in Chapter 6.4 require a total of 17,906,250 optimisation predictions[3] in total, which would make the simulation extremely computationally costly if full optimisation were required for each sub-problem.

[3]This value is in expectation, as the absolute number is stochastic.

6.5.3 Breaking Down Retrofit Decision Stages

The retrofit decision model presented so far, in the form of the ABM, is made up of three distinct stages. Initially, the household is triggered to consider a retrofit, modelled by a simple stochastic evaluation rate in Section 6.4. After being triggered, the household then needs to determine the nature of the retrofit that best suits their physical property and their household preferences. This optimisation problem has been delegated to the SO to provide a computationally cheap prediction of a WHRS given these properties in lieu of performing a full optimisation procedure, which does not scale when evaluating a large number of households and dwellings. The final element in the retrofit decision procedure is whether the agent installs the WHRS produced by the model, which will occur if the performance of the model, evaluated by the SEPM, financial and emissions models results in positive net utility given the household's environmental preferences. These stages can be summarised as:

1. Trigger Stage: The household is triggered to consider a WHRS. Up to this point, this occurs stochastically with a fixed probability based on an exogenously provided base rate of retrofit evaluation.
2. Selection Stage: The household discovered the best WHRSs based on their dwelling's properties and their environmental preferences. This is predicted by the SO.
3. Evaluation Stage: The household evaluates the best WHRSs to determine if the resultant NPV and emissions savings yield a positive utility based on their environmental preferences. The SEPM, financial cost, financial savings, emissions cost and emissions savings models are used to obtain objective values, which are evaluated with the household's utility function. Positive utility retrofit solutions are installed.

These three stages make up the decision model as used in this chapter. While the selection and evaluation stages have been the main focus of the models constructed thus far, the deconstruction of the decision process into these stages gives additional power to the modeller to add levels of detail for individual agents. The evaluation stage, for example, could add additional constraints that prevent the household from installing a retrofit even if it yields a positive utility. These could be based on the availability of necessary finance, their levels of risk aversion or subjective components such as the level of perceived disruption that installing the retrofit may bear.

While these post-selection alterations at the evaluation stage yield interesting potential for future work, the trigger stage is also of interest. The use of a base rate for triggering retrofit evaluation is an example of a top-down modelling approach, as the observation of an aggregated trend is being applied to individuals with no power to distinguish between their heterogeneous attributes. Chapter 7 will focus on extending the retrofit trigger model to encompass a greater level of detail using survey responses and individual responses to different triggers.

Chapter 7

Retrofit Trigger Modelling

> I have the simplest tastes. I am always satisfied with the best.
>
> Oscar Wilde

The retrofit adoption models presented in Chapters 5 and 6 can be broken down into three discrete stages. The first stage is for the household to be triggered to evaluate the possibility of installing a retrofit. In the prior model, a simple statistical trigger model was used based on the probability that a household will consider their retrofit potential in a given year. The second stage, the selection of the optimal retrofit for a household's preferences, was achieved with the use of the Surrogate Optimisation procedure laid out in the previous chapters. The final stage, evaluation of the optimal retrofit to determine if it is preferable to the status quo, was performed with the use of the Surrogate Energy Performance Model combined with the financial and emissions models described in Sections 5.2, 5.3.1, and 6.1.2 respectively. This chapter will investigate the replacement of the simple homogeneous trigger model of the first stage with a data driven model which accounts for heterogeneous agent responses to possible external retrofit triggers. The model will be designed and calibrated using data from a pilot survey conducted during the period of research.

NOTE: The pilot study presented in this chapter was conducted as part of a larger project, with members of that project making key decisions regarding the experimental design. While I was consulted prior to the experiment design phase, the majority of this was carried out by other members of the team. My contribution to this research was primarily in the processing and application of the results.

7.1 Survey Design

7.1.1 Best Worst Scaling

Best worst scaling (BWS) is a methodology which provides a list of options to participants who are asked to select the best and the worst option. There are typically three types (or 'cases') of BWS labelled simple Case 1, Case 2 and Case 3 [212]. Case 1 is the simplest case and was used in the design of the survey analysed in this section. Case 1 survey design provides high level 'object' style questions, relying on the participant to select an option based on their holistic understanding of the properties of the object at hand. This is in contrast to Case 2 and 3 which ask the participant to rank attributes of the object from which the experimenter can determine the underlying causes of given preferences. An example of a Case 1 question may be to ask participants whether they prefer to install a solar panel or additional insulation. A Case 2 version of this question may ask participants if they prefer to generate electricity to be used in their home or reduce the total amount of electricity required. Surveys can be extended to Case 3 by providing additional attributes within each question, for example, the cost and disruption levels associated with each individual installation.

The use of Case 1 BWS design in this survey had some advantages. In particular, it greatly decreases the number of questions required of participants, requiring a minimum of one question per object class [213, 214]. This keeps the time required to conduct the survey low and increases the number of topics and areas that can be studied. In this instance, only two questions from the survey were required to cast insight into the types of information that influence potential decision makers. The downside of this lower fidelity approach, however, is the detail of the data driven models that can be derived from them. With only two data points per participant per question, it is not possible to derive the entire preference set for a given individual. It is also not possible to determine which characteristics of the object drive a given choice, unlike in the Case 2 and 3 designs. Nonetheless, it has been empirically demonstrated (although not formally proven) that the best minus worst scores are linearly related to the fully constructed conditional logit models constructed from more complete data sets [213].

7.1.2 Pilot Survey Implementation

The pilot survey was conducted under the grant title "RCUK Innovation Fellowship in UK Housing Stock Decarbonisation" (Ref ES/S001670/1). This grant was sponsored by the Economic and Social Research Council (ESRC)

The survey was commissioned by Energy Systems Catapult[1], acting as sub-contractor to the University of Sheffield. This was compliant with the ethics board policy of the University of Sheffield by using an authorised sub-contractor who provided robustly anonymised data. No identifiable data was handled by members of the research group.

7.1.3 Survey Questions

Two questions relevant to this research were posed during the pilot survey. The first question related to the significance of possible information sources which could influence participants. The second question is related to factors that may trigger the participant to consider adopting a WHRS. The question was worded as follows for both:

> When thinking about investing in a home improvement, which of the following things would have the MOST and the LEAST influence on your decision. Please select one answer for MOST and one answer for LEAST.

The possible information sources are shown in Table 7.1 and the possible sources of triggers can be seen in Table 7.2. Due to the long form descriptions used in these questions, the visualisations in this chapter may use the answer code (indicated in the tables references) as a stand-in for the full description.

As well as the questions of interest, demographic information was collected regarding participants. Demographic information was collected regarding the following attributes of the participants, their current household, and property:

- Property type footnote
- Property age

[1]Address: 7th Floor, Cannon House, The Priory Queensway, Birmingham, B4 6BS E-mail: procurement@es.catapult.org.uk Telephone: +44 121 203 3700. Legal: Ian Jones / Ian.Jones@es.catapult.org.uk.

Code	Description
1	An expert from energy company or social organisation
2	Colleagues (e.g. work, school, college)
3	Entertainment media (e.g. radio, tv programmes, documentaries)
4	Family
5	Friends
6	Internet ads
7	Mailbox, leaflets, and brochures
8	Neighbours
9	News (tv, radio, online, or newspaper)
10	Social media
11	Street ads (e.g. billboards, bus shelters, and phoneboxes)

Table 7.1: Information source BWS choices and their corresponding codes.

Code	Description
1	Building a property extension
2	Your boiler breaks down
3	Full renovation/New “fixer upper” project – before you move in
4	When gas boilers become obsolete
5	A massive discount / clearance sale

Table 7.2: Trigger BWS choices and their corresponding codes.

- Property heating method
- Number of bedrooms in the property
- Household size (number of people who live at the property)
- The presence of anyone under 70 living in the dwelling
- The presence of anyone over 70 living in the dwelling
- The maximum level of education undertaken by the participant
- Year of the participant’s birth
- Whether the household’s disposable income exceeds £900 per month

These demographic details were chosen not just to comply with participant anonymity requirements but also for their direct mapping to fields present in the English Housing

Survey (EHS). Given the EHS was used as part of the disaggregated data set discussed in Section 4.5, many of these fields directly mirror those available for the agents in our simulated version of Nottingham.

7.1.4 Survey Sample

The pilot survey was answered by 1,138 online participants. Examining the responses of participants to demographic questions can demonstrate see that the sample of survey respondents is generally similar to the whole city data set used in previous chapters. This is shown in Table 7.3. It should be noted that some missing values were present in the survey results while the disaggregation method removed these from the Nottingham buildings data set, which may result in higher proportions in certain city data set characteristics as these percentages will sum to 100.

Attribute	Attribute Value	Survey Respondents	City Data Set
Household Size	1 member	4.0%	5.1%
	2 members	11.9%	14.2%
	3 members	24.3%	28.9%
	4 members	30.1%	28.2%
	5 members	19.7%	17.3%
	6 or more members	7.5%	6.3%
Property Type	Semi-Detached	19.3%	31.5%
	Mid-Terrace	15.7%	27.1%
	End of Terrace	13.5%	17.5%
	Detached	12.9%	13.9%
	Flat	33.1%	0.0%
	Bungalow	4.5%	7.6%
	Other	1.0%	2.4%
Under 16s?	TRUE	20.1%	24.3%
	FALSE	74.5%	75.7%
Over 70s?	TRUE	27.7%	25.9%
	FALSE	70.1%	74.1%
Education	Higher (e.g. Bachelors Degree)	43.5%	42.9%
	Sixth Form/College (e.g. A-levels)	24.6%	26.1%
	Secondary school (e.g. GCSEs)	17.7%	18.1%
	Postgraduate (e.g. Masters)	11.3%	11.0%
	Primary school	1.1%	1.4%
	No formal schooling	0.4%	0.5%
Bedrooms	1	3.7%	NA
	2	15.3%	NA
	3	38.0%	NA
	4	32.2%	NA
	5	9.5%	NA
Disposable Income	Less than £900	50%	NA
	More than £900	49%	NA

Table 7.3: Summary of demographic data shows survey was generally similar to city data.

7.2 Data Exploration

7.2.1 Best Worst Scaling Metrics

In order to evaluate and explore the BWS data visually, there are certain metrics that we can use to grant good interpretations. Initially, an examination of the raw best and worst counts for each response can provide a general understanding of the responses to different options. The next metric, best minus worst, simply subtracts the number of times an object was selected as the worst from the number of times it was selected as the best. This value is useful in providing a simple single value for each object and has been found to correlate highly with more sophisticated measures of each preference such as the results of conditional logit models applied to each question [213]. The downside of this approach is that it can be hard to distinguish choices that are highly controversial (and therefore have a large number of likes and dislikes) from values that are simply considered insignificant and therefore have a small number of both. Another useful BWS metric is best over worst, the ratio of best to worst selections for each choice option. While still susceptible to the same potential scaling effect, these ratios can give a good understanding of scenarios where a strong preference for one option over the other is present, even in cases where the absolute number of choices was not large. The final metric we will consider is relative importance. This value takes the difference between best and worst responses and normalises them to the number of responses of the most popular choice. This results in a value between 0 and 1 of the significance of the object at hand, proportional to the most preferred option. This has the benefit of being linearly related to the utilities calculated by a full preference score model [213] as well as taking ratio values between 0 and 1, which will be useful properties for the construction of the trigger model in Section 7.3.1.

Some other techniques to analyse the BWS results include clustering, a form of unsupervised ML which attempts to discover similarities between respondents which may provide insight into underlying patterns in the data. This has the potential for discovering additional archetypes by finding similarity across features that may not be obvious even to domain experts. An example of using clustering on BWS data can be seen in Auger et al. (2007) [215]. Another method of analysing results involves the construction of a conditional logit model, which can be used to find utility value coefficients for each possible choice. However, this generally requires more than a single question for each participant, relying on a series of questions to determine the ordered rankings of those options not selected as the single best or worst from within the set.

7.2.2 Information Source Responses

The key BWS metrics relating to the significance of different information sources can be seen in Figure 7.1. To interpret the option code in the diagram, refer back to Table 7.1 which contains the full descriptions as they were provided to participants. The results demonstrate that the most significant information source is perceived to be expert opinion. This result is robust among all the metrics considered. The least reliable source of information was perceived to be internet ads, with other forms of ads, such as leaflets or street ads also being viewed negatively. Opinions of family and colleagues were the highest rated social sources of information, which should be considered by any modeller looking to construct an influence model based on social network interactions. Friends and neighbours, on the other hand, were not seen as high quality information sources regarding WHRSs. The only other choice for which 'best' answers outnumbered 'least' answers included entertainment media, with more participants rating this source (including radio, tv programmes, and documentaries) as higher rather than lower.

7.2.3 Triggers Responses

The key BWS metrics relating to the significance of different WHRS evaluation triggers can be seen in Figure 7.1. To interpret the option code in the diagram, refer back to Table 7.2 which contains the full descriptions as they were provided to participants. The most significant trigger by all metrics was the adding of a property extension, with over 300 respondents recording this as the most significant trigger compared with less than 150 marking it the least. Opinions were more divided on triggers than with information sources, yet, boiler breakdowns received a significant but nearly equal number of best and worst responses. The least significant triggers were perceived to be large discounts and gas boiler obsolescence, although with relative importance scores of 0.53 and 0.47 respectively, there is still a reasonable level of significance applied to these triggers. These values will also be broken down by household archetype in Section 7.3.1, which will show that the preference order of triggers differs among the different categories of respondents.

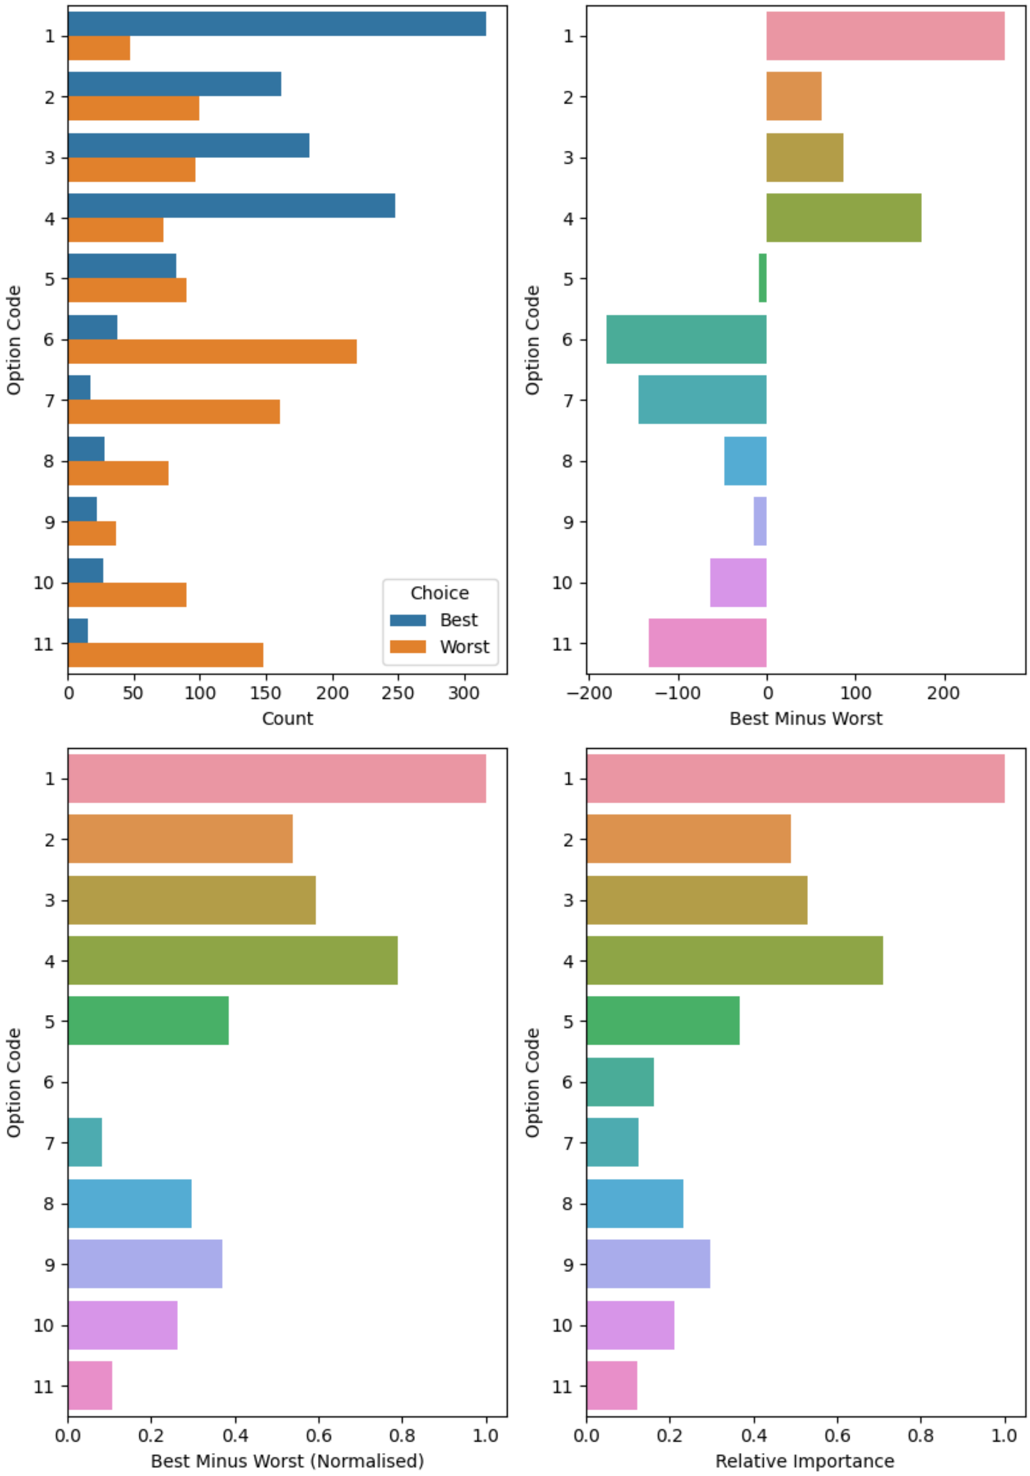

Figure 7.1: Key BWS metrics relating to the information source question.

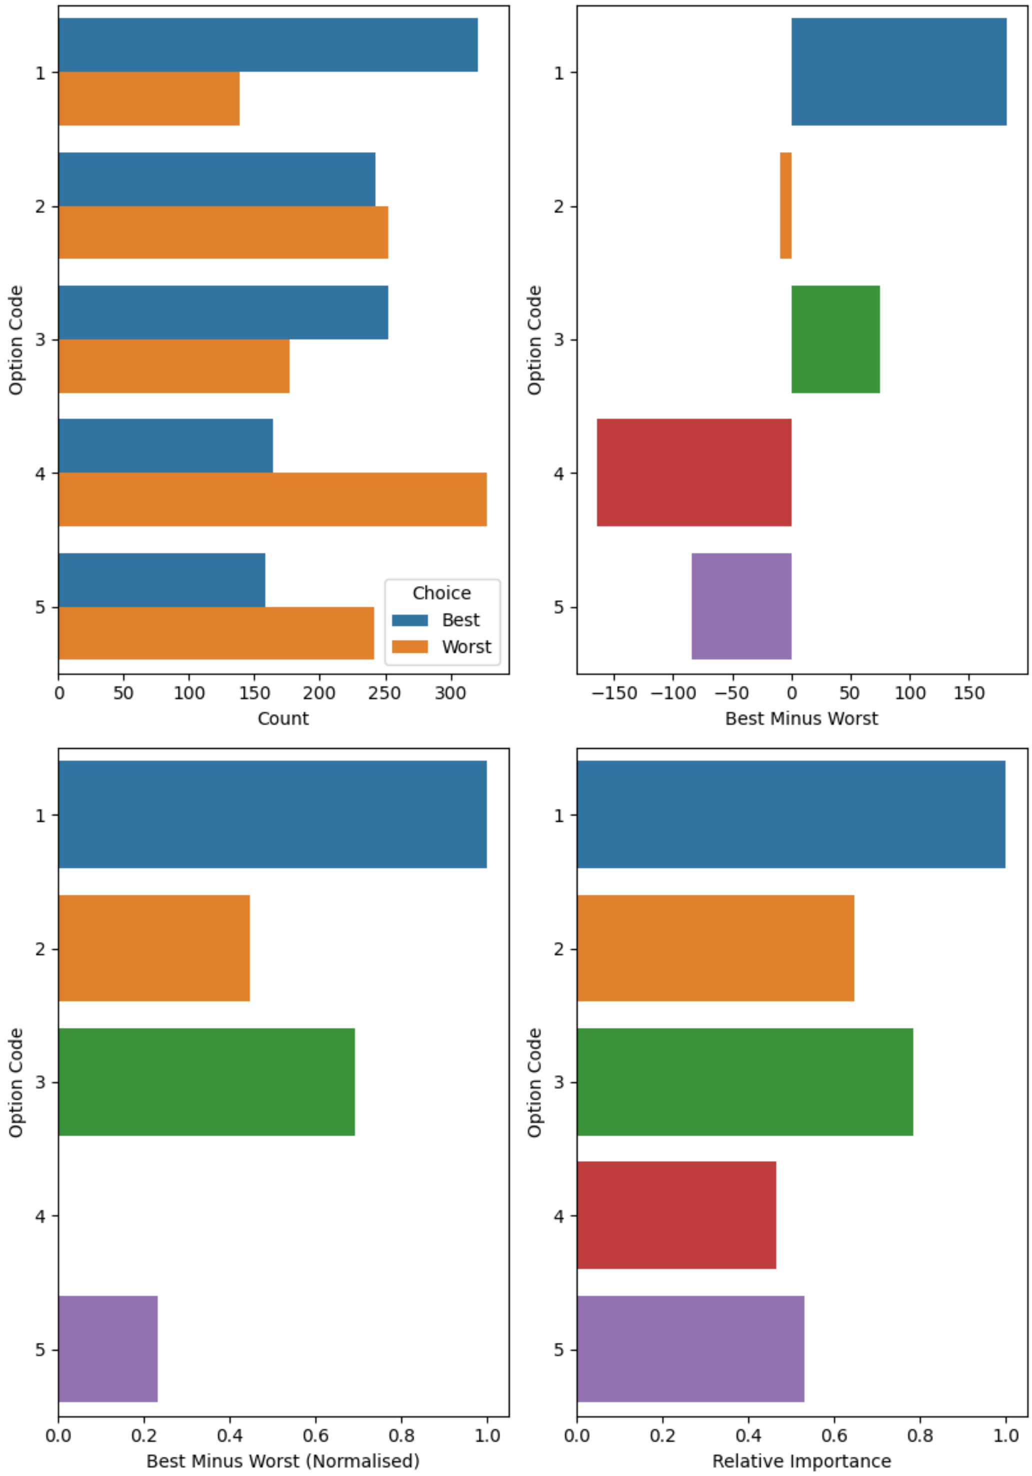

Figure 7.2: Key BWS metrics relating to trigger questions.

7.3 Designing a Trigger Model with Best Worst Scaling Data

7.3.1 Trigger Model Design

In Section 7.2 the relative importance of different potential retrofit evaluation triggers was calculated. Using a base evaluation rate similar to that used in the ABM design of Sections 5.5 and 6.4, it is possible to create a probabilistic model which is tailored to the stated preferences from our BWS results.

Consider first that if there was only a single trigger, the probability of a retrofit evaluation occurring, labelled as P(r), is the product of the probability of that trigger occurring, P(t), and the probability of a retrofit evaluation given that trigger.

$$P(r) = P(r|t) * P(t) \tag{7.1}$$

When this is extended to multiple triggers, the probability of any given trigger occurring ($P(t_i)$) is the complement of the probability of no events triggering a retrofit:

$$P(r) = 1 - \prod 1 - (P(r|t_i) * P(t_i)) \tag{7.2}$$

While the absolute values of $P(r|t_i)$ cannot be gleaned from the BWS, we do know the relative importance of each trigger to survey participants using the scaled best minus worst values. If we consider the most significant trigger, t_0, and the rest of the triggers to have a relative importance score of σ_i then the probability of participants considering a retrofit can be expressed as a linear combination of the most significant factor:

$$P(r|t_i) = \sigma_i P(r|t_0) \tag{7.3}$$

Substituting these values into Equation 7.3.1 allows the expression of the base retrofit probability in terms of only the probability of each trigger and the probability of installing a retrofit given the most significant trigger occurs. If the base retrofit probability and the probability of each trigger are taken as inputs, this allows us to solve for the remaining values: the probability of installing a retrofit given a specific trigger has occurred.

$$P(r) = 1 - \prod (1 - \sigma_i * (P(r|t_0) * P(t_i))) \quad (7.4)$$

7.3.2 Trigger Model Training

Initially, we will use the entire BWS data set to train a retrofit evaluation trigger model. This will result in a homogeneous agent design, with all agents reacting equally to a given trigger. In order to calibrate the conditional probabilities of the proposed triggers for the model design proposed above two main sets of parameters are required. Initially, the base rate of retrofit installations is required. This is the number of retrofits evaluations expected to occur in a given year independent of what triggers those retrofit evaluations. Luckily this is a parameter we have worked with before and represents the evaluation rate discussed in Sections 5.5 and 6.4. The same base evaluation rate of 1% will be used in this section for consistency with these previous sections. This also ensures that the underlying number of retrofits will, at present, remain unchanged. The remaining parameters, which are shown in Table 7.5, are the probabilities these events occur in a given year.

The selection of the annualised probabilities of given triggers is an exercise in data collection combined with model conceptualisation. The estimated probability of building an extension was based on the 2021 planning permission data provided by Nottingham City Council [216]. The average probability of a boiler breaking down was calculated based on the annualised expected lifecycle of a boiler, adjusted to account for the proportion of dwellings in which a boiler is found. The probability of a renovation before moving in was based on the average rate of relocation rate of approximately 7% per year in England among homeowners [217]. The final base probabilities required some additional conceptualisation. Given that gas boilers have no yet become obsolete, this event is not contributing to the base rate of installations that contribute to the model. As such, the probability has been set to zero. This will still allow the effect on retrofit probability to

Code	Trigger	P(Trigger)	Relative Importance	P(Retrofit|Trigger)
1	Building an extension	0.02	1.000	0.099
2	Boiler breaks down	0.04	0.646	0.064
3	Renovation before moving in	0.07	0.785	0.078
4	Gas boiler obsolescence	0.00	0.465	0.046
5	Large discount	0.00	0.532	0.053

Table 7.4: Trigger model for homogeneous agents with a base retrofit evaluation rate of 0.01.

be calculated, as the relative importance of the event, should it occur, has been ascertained from the BWS responses. In this case, the conditional probability should be that if households begin to perceive gas boilers as becoming obsolete, how likely would they be to consider a WHRS evaluation. A similar approach has been taken with the 'large discount' trigger. It has been conceptualised as a particularly large discount (possibly driven by government subsidies, tax cuts, or changes to the market) that are not in place when the base rate occurs. This allows for the consideration of a discount that occurs as an exogenous policy position in the ABM. It should be noted though that a discount would affect the optimisation parameters under consideration and require the retraining of the SO in order to evaluate the optimal retrofit installations under the new financial model.

The outcome of the trigger model's training can be seen in Section 7.4. The P(Retrofit|Trigger) columns are the calculated conditional probabilities that can be interpreted as the probability that a household will evaluate retrofit potential given that a specific trigger has occurred. For example, the most significant trigger, building an extension, will lead to an evaluation of retrofit potential 9.9% of the time, while a household perceiving gas boilers to be obsolete will only lead to a retrofit evaluation in 4.6% of cases.

7.3.3 Modelling Heterogeneous Agents Using Household Archetypes

The base trigger model discussed above makes the strong assumption that retrofit triggers affect all agents equally. It is possible to stratify the BWS data to construct a heterogeneous trigger model by breaking respondents down into household archetypes. Given the availability of data across both the Nottingham city housing data set and the BWS demographic data, it is possible to break households down into stereotypes that depend on both the age of the primary household breadwinners and the presence of dependent children. Splitting the BWS data into these archetypes produces decision model parameters shown in Table 7.5.

Trigger	Household Archetype	Rel. Importance	P(Retrofit|Trigger)
1	couple, no dependent child(ren) under 60	1.000	0.086
2	couple, no dependent child(ren) under 60	0.848	0.073
3	couple, no dependent child(ren) under 60	0.904	0.077
4	couple, no dependent child(ren) under 60	0.411	0.035
5	couple, no dependent child(ren) under 60	0.580	0.050
1	lone parent with dependent child(ren)	0.987	0.090
2	lone parent with dependent child(ren)	1.000	0.091
3	lone parent with dependent child(ren)	0.725	0.066
4	lone parent with dependent child(ren)	0.517	0.047
5	lone parent with dependent child(ren)	0.957	0.087
1	couple with dependent child(ren)	1.000	0.106
2	couple with dependent child(ren)	0.563	0.060
3	couple with dependent child(ren)	0.740	0.079
4	couple with dependent child(ren)	0.454	0.048
5	couple with dependent child(ren)	0.462	0.049
1	other	0.577	0.061
2	other	0.333	0.035
3	other	1.000	0.106
4	other	0.183	0.019
5	other	0.333	0.035
1	couple, no dependent child(ren) aged 60 or over	0.452	0.057
2	couple, no dependent child(ren) aged 60 or over	1.000	0.126
3	couple, no dependent child(ren) aged 60 or over	0.436	0.055
4	couple, no dependent child(ren) aged 60 or over	0.350	0.044
5	couple, no dependent child(ren) aged 60 or over	0.655	0.083

Table 7.5: Trigger model for heterogeneous household archetypes.

As can be seen from Table 7.5, the different archetypes manifest significant differences in the relative importance each one places on given triggers. This results in significant deviations in the probability of a retrofit evaluation when a trigger occurs. Couples aged over 60, for example, are significantly less likely to consider a retrofit installation when building an extension with a conditional probability of only 5.7% compared with 9.9% in the homogeneous trigger model above. Lone parents with dependent children appear to be the most price sensitive demographic as they are the most influenced by the possibility of a large discount/sale with a 0.957 relative importance. Differences in these calculated conditional probabilities will result in disparate outcomes across archetypes when underlying circumstances change within the ABM.

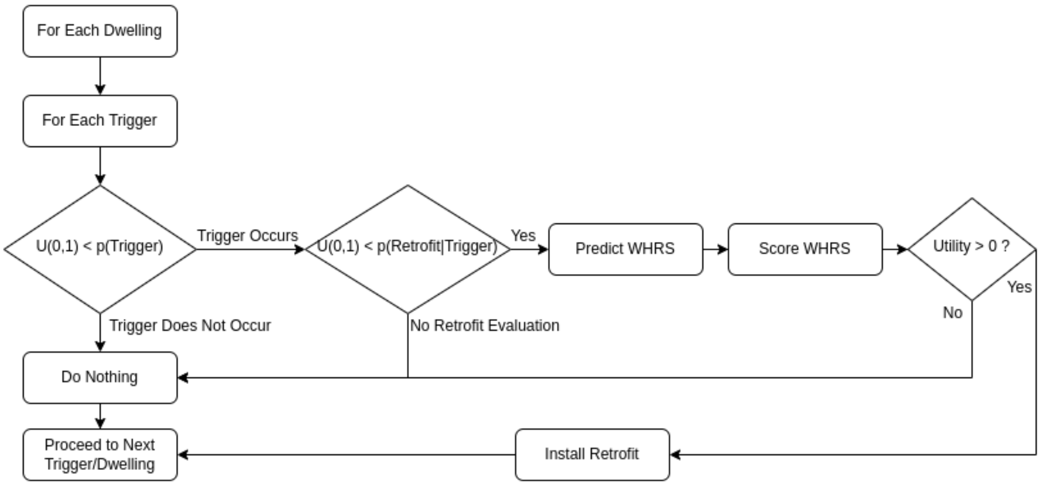

Figure 7.3: Flow chart of the retrofit decision model using the calibrated trigger model.

7.3.4 Integrating the Archetype Trigger Model into an ABM

Given the overlap between the demographic data collected during the BWS survey and the EHS data used to generate the Nottingham buildings' data set used for the ABMs up to this point, it is straightforward to map household archetypes onto the decision model archetypes discussed above. The archetype-specific trigger models described above can then be used to observe how households will respond to different changes in the prevalence of retrofit triggers.

A flow chart showing how the trigger model described integrates into the ABM can be seen in Figure 7.3. As the trigger model is modular, much of the decision model used in Section 6.4 is unchanged. The method of predicting WHRSs using the integrated SO as well as the scoring procedure of retrofits has remained the same. However, the initial decision to evaluate is now triggered based on the underlying probability that the triggers occur, combined with that agent's conditional probability of evaluating retrofit potential once the trigger occurs. This adds additional heterogeneity to the ABM, as now not only is each dwelling being uniquely modelled, but the agents also have different propensity to perform retrofit evaluations based on their archetype.

7.3.4.a Modelling Perceived Obsolescence

With the introduction of heterogeneous household trigger models it is possible to model the impact of certain scenarios that would not have been possible in earlier iterations of the ABM presented in Sections 5.5 and 6.4. Consider, for example, the obsolescence trigger, which occurs when a given household perceives their gas boiler to have become obsolete. With the heterogeneous trigger model in place, it is possible to model the impact of different obsolescence scenarios on both the overall uptake of retrofits as well as a breakdown of how each demographic is impacted.

Given the rate of perceived gas boiler obsolescence is unknown, a set of scenarios using a simple linear obsolescence model will be used for demonstrative purposes. The model will apply an obsolescence rate which leads to increasing perceived obsolescence each year. The low obsolescence scenario will assume an obsolescence rate of 0.5% per year, while the rapid obsolescence model will assume a rate of 1%. The experiment was repeated 5 times for each scenario. No government campaigns affecting the evaluation rate were simulated to ensure that the effects seen were a result of the obsolescence model's effect on the emergent evaluation rate. Environmental awareness campaigns, as discussed in Section 6.4, were included to raise HCVs after their decision to install a retrofit had been triggered. The penetration rate of these campaigns was set at 10%.

The results of the obsolescence scenarios can be seen in Figure 7.4, which shows the impact of the different obsolescence models on different household archetypes. Results are presented in terms of retrofit rates (the number of retrofits per dwelling per year) to control for the different provenances of household archetypes. An initial feature of interest is that the baseline retrofit rates differ between archetypes despite the base evaluation rate used for calibrating their trigger model does not differ. This is caused by differences in the initial states of the dwellings, resulting in different provenances of positive utility retrofit across archetypes. Investigating the effect of the obsolescence models, it can be seen that the largest impact of higher obsolescence scenarios is on lone parents with dependent children. Among this archetype, the fast obsolescence scenario resulted in almost twice as many annual retrofit installations over the time period compared with the baseline. In contrast, the group making up the 'other' archetype faced the least impact from perceived obsolescence, an expected result given the low relative importance score of 0.183 for this retrofit amongst the BWS sample. It is notable that even these conservative estimates for perceived obsolescence rates result in significant boosts in retrofit installation rates.

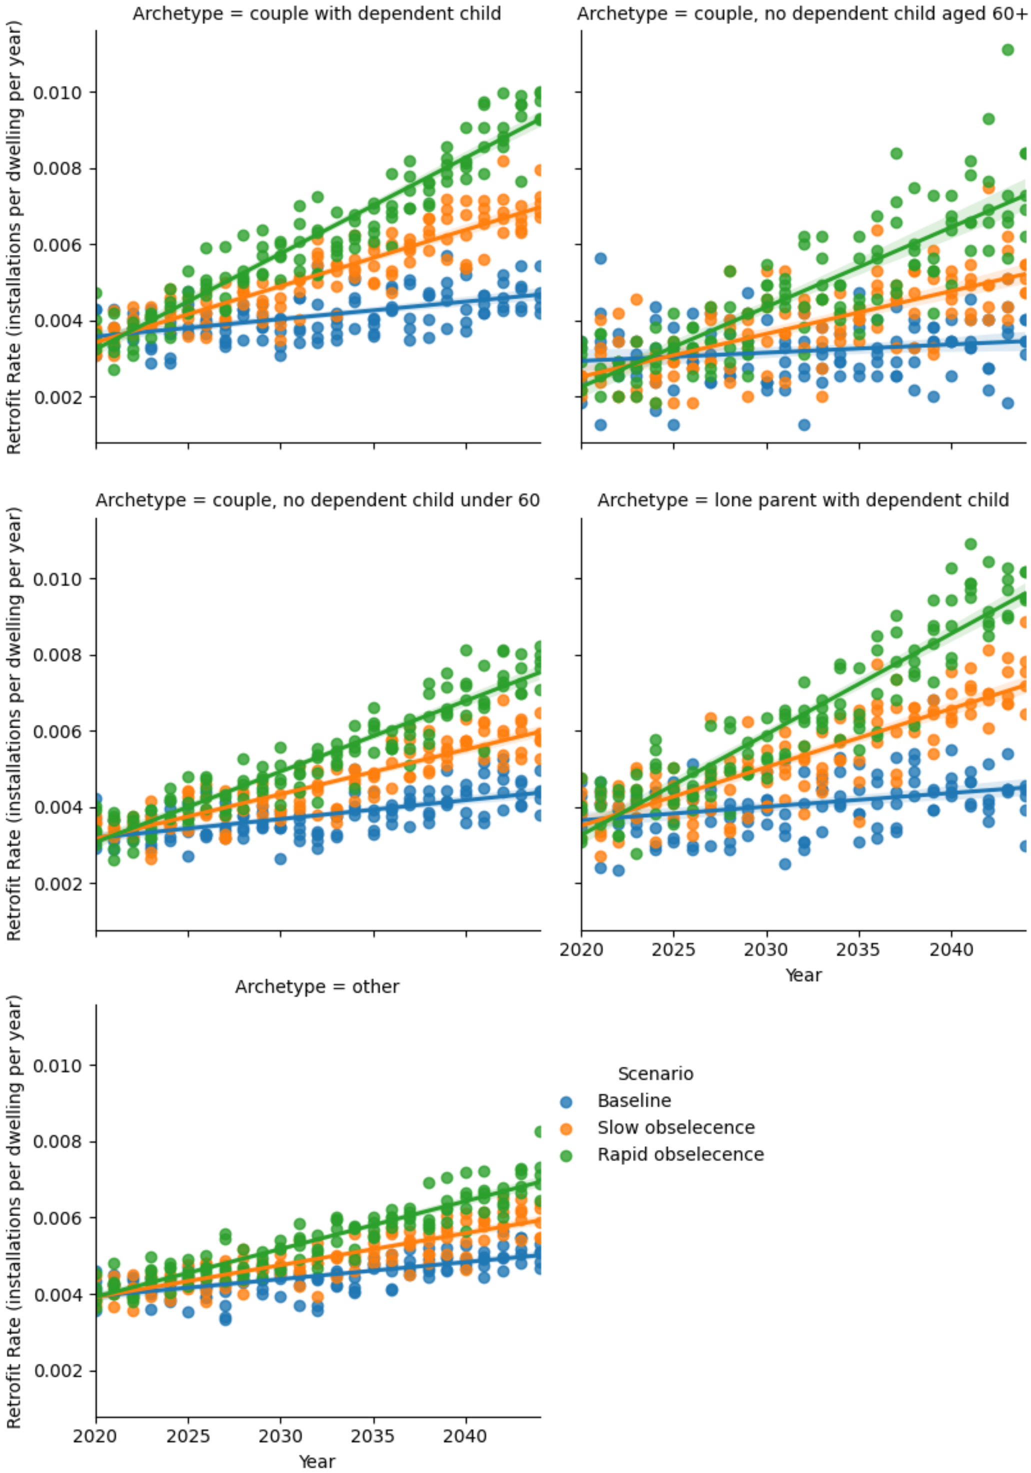

Figure 7.4: Increases in perceived obsolescence resulted in different effect sizes across household archetypes.

7.4 Discussion

7.4.1 Integration of information source data into ABM

While the survey responses relating to information sources was analysed in this model, the trigger model constructed and integrated into the ABM was reliant only on the BWS data relating to potential retrofit triggers. Constructing a social influence model to evaluate the level of environmentalism among households is out of scope for this research and would likely require significantly more data collection phases. However, the information source answers in 7.2 do give some indication as to what such a model would look like as well as what the most significant components would be. In order to integrate this with the SO constructed in Section 6.3, the target feature of an environmentalism model would be the WTP per ton of emissions mitigation. In order to obtain data using this survey methodology, a Case 2 or Case 3 BWS model would be required. These methods breakdown the features of given retrofit choices, such as the cost, energy saving, and emissions mitigated and allow the modeller to calculate the marginal utility associated with each unit of each attribute. Doing so would enable the construction of a utility function of the form discussed in Chapter 6 on an individual level for each participant. Clustering could then be used to assign utility values to the simulated households based on their similarity to members of the surveyed group.

As well as a static utility model, the influence of social pressures could be integrated into an information source model. The data presented in this chapter indicates that more value is placed on the opinions of experts, colleagues, and family with less being placed on forms of paid advertising, neighbours, and friends. These values could be used for the calibration of a dynamic social model which captures these interactions.

7.4.2 Relaxation of trigger model assumptions

The trigger model discussed in this section is reliant on some assumptions discussed during its design. It relies, for example, on a known base rate of retrofit evaluation triggers in order to calibrate the trigger model. Additionally, triggers are modelled as random and independently occurring events with a fixed distribution. The triggers discussed in the survey were also considered the dominant causes of retrofit evaluation, with the excluded triggers assumed to contribute minimally to the total retrofit evaluation rates. It is worth considering the conditions under which these assumptions could be removed in future work or with more data availability.

A key assumption in the trigger model design was that of a known base rate of retrofit. This was required in order to calibrate the model using the relative importance of the different triggers obtained from the BWS data. This assumption could be relaxed with the introduction of a more thorough BWS survey, which would allow the absolute utility coefficients to be calculated, rather than relying on relative values [218]. This would allow the modeller to derive the retrofit evaluation rate endogenously based on the trigger frequencies rather than relying on an exogenous rate used for model calibration. Such a survey design would require asking respondents significantly more questions, as well as gathering more significant demographic information to ensure responses can be calculated at an individual rather than archetypal resolution.

Another assumption used in implementing the trigger model was that the prevalence of potential triggers were both independent and could be modelled as simple distributions. While there is no significant methodological issue with this, the relaxation of the assumption could grant additional control to modellers to construct more granular models. Consider for example the trigger which occurs when a household moves dwelling. In this chapter, moving was modelled as an independent random variable occurring at a fixed annual probability. Modellers with a particular interest in the relationship between the housing market and the adoption of WHRSs could instead use a more sophisticated dwelling-moving model, such as an agent-based approach in which all households consider moving based on observations such as their personal finances, the job market, house prices, and interest rates. The model is agnostic to the cause of the move, making this extension straightforward to add once the relocation model has been constructed.

A final assumption that is worth considering is the presence of other, unconsidered triggers in the model. One of the restrictive aspects of any BWS question design is the requirement for options to fit into predetermined categories selected by the modeller. In the design of the survey in this chapter, potential triggers were selected after surveying the literature as well as consulting with industry experts to ensure that the most significant triggers were added. However, there is still relatively little understanding of all the factors that drive domestic WHRS-adoption, leaving the potential for missing triggers, and neglecting any decisions that do not have an obvious or relevant trigger. This problem is confounded by the self-reporting nature of BWS surveys, as a trigger that is perceived as significant by an individual may not actually be the most significant when faced with real-world decisions [219].

Chapter 8

Discussion and Conclusion

8.1 Summary and Discussion of Conducted Work

The main chapters relating to the contributions made in this work, Chapters 5, 6 and 7, were accompanied with discussion sections relating to the specific concepts introduced within those chapters, including some limitations and potential for future work. This section will attempt to summarise these, as well as add some general points of discussion not included in those sections. This is because these further points of discussion apply more generally than the scope of the given chapter, such as an overview of this research's contribution, or they relied upon context that was not in place until later chapters.

8.1.1 Single-Objective Surrogate Optimisation

Chapter 5 explored the framing of agents as rational and self-interested actors. This allowed for a utility function made up of a single objective, Net Present Value (NPV), which represented the lifecycle financial returns of the model. The single-objective retrofit optimisation was performed with the prior intention to extend the process to two objectives at a later stage but it required many preliminary steps of conceptual modelling and data collection which would extend to the full multiple objectives in later chapters.

The preliminary stages required to construct the single-objective SO were quite significant. A process for constructing energy simulation models from the available data set resulted in a distribution of simulated energy demand values matching those recorded in the real data set. The SEPM constructed from the simulated energy demand data set was of very good quality, with small and normally distributed errors. This SEPM was then used alongside a financial model to predict the NPV of any given WHRS, a form of evaluation that could be integrated into a GA to obtain a large and representative set of WHRSs. This set of WHRSs was used to construct a single-objective SO, using a combination of classification and regression neural network models. The single-objective SO constructed in this work performed well, with well-performing classification and regression metrics that reflected a reasonable level of accuracy. The SO was able to identify the majority of WHRSs within the stock, with only a small reduction in the mean NPV compared with the test sample. Due to the speed of the SO compared with the traditional optimisation procedure, it was possible to run a stochastic ABM with many repetitions without significant computational time or cost.

This work was significant in laying the preliminary modelling stages necessary to expand the technique to include environmentally-conscious agents. Many of the conceptualisations, assumptions, and simplifications made at this stage were carried through to later stages. While later work was focused on the non-financial aspects of WHRS-adoption, it is important to note that financial and energy savings are key factors in retrofit decisions, so a single-factor analysis focused on these components was deemed important. The purely rational model is also a good point of comparison when evaluating the multi-objective problem later, as these agents can be considered a special case of environmental agents who have a carbon valuation of 0£/tCO2e. Whilst the agent decision models used at this stage were more rigid in their assumptions than in later chapters, the trained SO provided value independently of those assumptions. The ability to identify which dwellings are most likely to have WHRSs of positive NPV with spacial resolution to the single dwelling is a powerful tool in and of itself. The speed with which the SO is able to predict these NPVs makes the problem highly scalable, limited only by the availability of feature data to make predictions for a given building. Given the data sources used, the process could be applied to anywhere with OS map data, which encompasses the entire UK.

The solutions found during this procedure were only found in a subset of the possible solution space. When considering wall insulation, for example, rational agents were only selecting internal wall insulation due to its lower financial cost when compared with external insulation. This reduced optimisation space is context-specific, however, and a change in the inputs to the financial cost model could result in inverting this outcome or result in mixed outcomes where the decision of external or internal insulation is more contextual. Indeed, when environmental agents were introduced, even with the same financial cost model, a mixed scenario occurred whereby the choice of internal vs external insulation placement was dependent on the HCV of the household.

Whilst the work on rational agents was considered in terms of single objectives, it is worth noting that NPV is really a composite objective that takes into account the lifecycle costs and savings provided by the retrofit. This process of taking multiple objectives and coercing them into a single value, scalarisation, allowed the problem to be approached as if it were a single objective. This is not entirely dissimilar to the utility function scalarisation used in Chapter 6, although the weightings between cost and savings were, after accounting for discounting, fixed. This means that although there were multiple objectives scalarised to create a single-objective function, all agents had identical utility functions and can therefore be considered homogeneous with respect to their decision models. The heterogeneity in this section of work was obtained purely through the physical and economic properties of the WHRSs applied to the unique dwellings, rather than from the environmental preferences of the agents.

8.1.2 Multi-Objective Surrogate Optimisation

Extending the method to account for multiple objectives allowed relaxation of the assumption that households were only interested in the financial return granted by a WHRS. Given the literature discussed in Section 3.3, it is clear that decision makers are conscious of the environmental benefits of their actions and are willing to make financial trade-offs to account for them. The utility model used to measure the impact of environmental preferences was chosen to be agnostic to the underlying cause of environmentalism. This was because there is conflicting evidence on what the actual cause of environmentally-conscious behaviour is: pure altruism, an instinct for self-preservation, or a 'warm glow' that is felt when people perceive their actions to be virtuous. Rather than attempting to account for these conflicting root causes, the model relied on the principle of Willingness to Pay (WTP) per ton of emissions mitigation. Using this metric, the emissions cost and savings models were constructed to evaluate the environmental impact of each WHRS.

The integration of an environmental objective function into the process flow required some adaptations at the optimisation and SO stages. Instead of using a multi-objective adaptation of the GA, utility functions were used to scalarise and locate individual solutions with known HCVs. This process kept the HCV associated with Pareto solutions transparent, allowing them to be used as input features when training the multi-objective SO. The trained SO could then be provided with the details of a dwelling alongside the HCV to predict a WHRS suitable for the physical and social properties of the solution. The multi-objective SO had reduced performance metrics compared its single-objective counterpart, but was still generally of good quality. The glazing prediction was typically the worst-performing. This is was unsurprising, as it was the component that the traditional optimisation process was the least able to reliably identify. Given this was used as the SO training set, it reflects a messier underlying system rather than an incorrectly constructed model. The multi-objective SO was then integrated into an ABM in order to model the effect of campaigns targeting household environmentalism.

A model predicting building-specific WHRSs at this scale, especially covering heterogeneous environmental preferences, would not have been possible with the existing methods discovered in the literature. At this stage, several aspects of the ABM were still quite rudimentary. While the SO was able to provide high quality predictions for which retrofit households with a given environmental preference would install, little effort was placed on determining which households would initially consider installing a retrofit. The trigger model introduced at this stage involved a simple probability that households would evaluate retrofits in a given year. This method of modelling triggers is not necessarily flawed, as ultimately it does imitate the behaviour of households in a general sense: every year some households, independent of the reason, will choose to consider a WHRS. This trigger model is, however, limited in resolution in two key ways. Firstly, using an aggregate probability makes it difficult to predict which households are more likely to consider WHRSs, failing to account for the fact that due to different underlying causes some household archetypes are likely to consider retrofits at different rates, or due to different causes. Secondly, the low resolution of this trigger model limits the capacity of modellers to run certain experiments relating to retrofit triggers. Consider, for example, an experiment that wished to predict the effect of a subsidy that reduced the cost of a given set of WHRS components. Such an experiment could be run by retraining the SO with new cost data to determine the influence of the sale on the quantity of WHRSs installed at a given evaluation rate, but the effect of the sale on the evaluation rate itself would need to be provided exogenously by the modeller.

Expanding the optimisation objective, and thus a whole front of solutions for each dwelling had the additional effect of diversifying the solution outcomes to encompass more of the possible solution space. However, there were some possible solutions, such as internal wall insulation made of EPS, which were still not found in any Pareto optimal WHRS. This raises the question as to whether the material should be removed from prior stages to simplify the optimisation process and reduce the number of unused features in models found later in the process. The decision was made not to do this, as the methodology described in Chapter 4 is designed as a computational workflow to be performed iteratively and with some independence from the input values at any given stage. In the same way that the distribution of found solutions expanded when the assumption of rationality was lifted, we may expect a portion of the feasible solution space to become relevant again should inputs of the cost or emissions model were to change the underlying problem.

8.1.3 Retrofit Trigger Model

In order to combat the limitations of a simplistic retrofit evaluation, the trigger model discussed in Chapter 7, laid out a more sophisticated approach. Using BWS data collected from a pilot survey, the relative importance of different retrofit triggers was calculated. A model was calibrated using these values, combined with an exogenous base rate of retrofit evaluations. This ensured that the overall evaluation rate would mirror that of prior simulations whilst freeing the modeller to perform experiments relating to the underlying triggers. The BWS and housing data were both split into matching archetypes to account for heterogeneous agent triggers. There were significant differences in the relative importance of different triggers among groups. The archetype-based trigger model was then integrated into the ABM design, which was simple, due to the modular nature of the ABM. A demonstrative experiment was conducted, in which households began to perceive their current gas boiler as obsolete at increasing rates, resulting in diverging WHRS-outcomes across groups that would not have been observed in previous iterations of the model. It was found that couples with dependent children were the most responsive demographic to boiler obsolescence scenarios.

At this stage, households differed in both their environmental preferences and their responses to different external triggers occurring, as well as having unique dwellings with differing environmental and economic retrofit properties. This level of detail could not have been achieved at this stage using methods found in the literature. Sample and upscale approaches would not be able to model individual agents at such a level of detail, while traditional agent-based approaches were not technically feasible at the urban scale.

8.2 Limitations and Future Work

8.2.1 Energy Simulation Scope

As with all research, the scope of the work conducted has been limited to ensure feasibility within the time and resources available. One example of scope limitation within this research was the range of WHRS components considered, with solutions focused on wall, roof, and window insulation, as well as fuel sources. The most significant scope exclusion is likely that of innovative or generative retrofits: photovoltaics, wind sources, and heat pumps. These were excluded partly due to the complexity of integrating the complex dynamic power supply within the EnergyPlus framework, such as the required energy storage or dynamic feed-in pricing. The scope of the research was also limited with respect to material type, reducing the list to the three most common retrofit materials, to reduce the complexity of the optimisation problem, as well as ensure good data availability. Nonetheless, the process is easily adaptable, allowing for the addition of arbitrary materials (including hypothetical materials not yet in production). New materials can be incorporated by adding the data of their physical properties into the simulation and SEPM model, and details of their material and installation costs to the financial model. To extend this to the multi-objective case, the emissions model would also require data on the lifecycle emissions of a given material.

Another limit of the research conducted is in the modelling of human behaviour within the dwelling. Household behaviour is modelled using a simple heating setpoint and schedule, which is sufficient to capture the energy demand of a dwelling under normal conditions. However, more sophisticated energy demand models account for real-time interactions between the householders and their built environment, accounting for behaviours such as opening windows, the use of blinds, and the changing of temperature setpoints based on perceived thermal comfort. While these behaviours have been observed to account for some variance in energy demand, they were infeasible to include when dealing with the scale of the problem addressed by this research. The necessary data to model such behaviours was not available and capturing such elements in the SEPM would require significant further research. Additionally, such behaviours are likely to make the most difference when measuring cooling energy demand. Cooling was also excluded from the scope of this research as domestic cooling in the UK is very uncommon and the additional complexity added to all stages of modelling would have only accounted for a small number of dwellings for which overheating is a problem.

Future work could extend the simulation scope with the introduction of more sophisticated physical and behavioural modelling techniques. The implementation of additional physical components such as heat pumps or battery storage is likely to be more of an engineering problem than a matter relating to further research. In contrast, the implementation of behavioural components into the model is likely to be more challenging, as capturing these behaviours for the SEPM and SO input features is likely to be a challenge. Nonetheless, it is likely to be a necessary step in modelling hotter climates where cooling demand, air flow, and overheating, represent significant parts of a household's energy use.

8.2.2 Stochastic and Risk Aware Modelling

One practical limitation of the methods used in this research was the use of deterministic models that did not account for uncertainty. Consider the financial cost model, for example, which takes as input the properties of a dwelling and a proposed WHRS and outputs the cost of installing the solution. The model outputs an absolute value based on the expected costs of labour and materials to perform the WHRS. In the real-world system, however, there would be a degree of uncertainty, resulting in a range of possible costs depending on random factors as well as unknown details about the dwelling. Consider, for example, a dwelling may have a previously unknown damp problem that must be resolved before internal wall insulation can be installed. This would influence the cost of installation but cannot be fully accounted for in the financial cost model. To some extent, this property could be considered desirable, as the household making the decision would also be unaware of the additional cost. In this sense, using a deterministic expected value is close to the decision metric available to households themselves. This does not account for risk-aversion, as particularly risk-averse households may rationally avoid positive NPV solutions. In other places, this component is accounted for, such as the discount rate applied to future returns, which will account for a degree of risk-aversion within households. While the inclusion of stochastic or risk-aware models could add additional resolution to the procedure, it is worth noting that without high quality data on the risk aversion levels of households, this inclusion would not necessarily change the aggregate findings of the model. In a sense, using simulated risk aversion would simply add noise around the expected value models which were implemented.

While the choice to use deterministic models was justified, there is potential for future work in investigating how stochastic or risk-aware models could be integrated into the framework. Of particular interest would be whether the aggregated results do indeed differ from a deterministic model based on expected values. Some modelling methods, particularly Bayesian approaches, can allow for non-deterministic or risk-aware modelling and could be substituted in place of the existing methods. However, the availability of data is likely to be a significant stumbling block, as there are no large scale real-world data sets relating to the distribution of costs and savings which could easily be used for model training and validation. Where data is available, it is often aggregated and presented in wide ranges, as was seen in the validation data laid out in Section 5.3.1.b.

8.2.3 Optimisation Validation

Another limitation of the research conducted was the ability to externally validate some of the later stages of research. Where the target system of a given model was an external phenomenon, such as the energy demand simulations or cost model, sources of external data were obtained to ensure the models sufficiently replicated the behaviour of the system under examination. However, some of the stages, such as the SEPM, were only validated against the previous stage. Conceptually, this is acceptable, as the target system of each of these higher-level models is only made up of the sub-models that make it up. Consider the optimisation stage, for example. While the optimisation procedure is made up of validated sub-models, lack of data availability prevents the validation of optimal WHRSs against real models. As such, the SO was only validated against the optimisation data that was used to train it. This validation against simulated data is necessary, and it was ensured that both validation and test data were separated from training data to ensure model performance was accurately reflected. Nonetheless, validation against simulated data only may hide incorrect assumptions of other forms of error introduced at these stages. It would therefore be a potential area of future work to attempt to validate the retrofit decision model against real-world decisions made. Should the required data never become available, a survey methodology attempting to validate against theoretical decisions could provide some additional validity.

8.2.4 Household Carbon Valuations

The addition of HCV into the retrofit decision model provided a transparent and measurable indicator of household environmental preferences which was backed both empirically and theoretically, as laid out in Section 3.3. However, the use of WTP-based measures does come with some drawbacks and limitations. While there is WTP data available was used to estimate HCVs assigned during the ABM, the estimates vary quite significantly across domains and applications. While WTP is measurable, estimating values for specific households based on their characteristics presents a challenge. In this study, a simple stochastic allocation of HCVs was used, with all household values being drawn from the same distribution. There is potential for future work to construct more sophisticated models for the prediction of households' WTP. Such research could case two or three BWS surveys that would allow for the calculation of utility function coefficients which could be paired with relevant socio-demographic information to allow archetype-based allocation of HCV similarly to how the trigger model was constructed in Chapter 7.

8.2.5 Objective Function Design

While Chapter 6 extended a household's objective function to include a measure of environmentalism and further captured the heterogeneous behaviours of agents by modelling different responsiveness to changes in external triggers (Chapter 7), the number of variables taken into account when evaluating retrofit decisions is still limited. There are other factors, such as the level of disruption caused during installations, which are not currently accounted for in the model. While these factors are likely to be significantly less important than the economic and environmental impacts of the WHRSs, their omission will account for some error. If we consider for example the bias for internal wall insulation among agents with low HCVs, this was driven by the lower cost of internal wall insulation. However, the installation of internal insulation is more disruptive than external insulation, and factoring this element in may result in more agents selecting external installations, even when faced with a higher cost. The omission of this factor was intended to keep the scope of the research feasible. Future work could focus on methods of accounting for these other objectives, using either an adapted objective function or by using a pre or post optimisation method such as a trigger model.

8.3 Concluding Remarks

Domestic dwellings account for over a third of the national energy demand and approximately a quarter of total CO_2 emissions in the UK, with the majority of this energy demand being used for electric or gas based space heating [220, 7]. While significant progress has been made in the design of new buildings for efficient heat generation and retention, most of the existing stock will still be occupied by the UK government's deadline of 2050 for net zero emissions [5]. Meeting this target will require strategies to transform the existing building stock with Whole House Retrofit Solutions (WHRSs). To analyse such policies, an increased emphasis should be placed on robust tools for modelling and optimising the policy decisions such as tax rate, public engagement, and methods of outreach. However, detailed dynamic analysis of building stock retrofit adoption is difficult, as not only do the physical properties of each dwelling differ, but each household has heterogeneous preferences in selecting a WHRS to meet their objectives.

Existing methods for modelling the adoption of WHRSs have several limitations. Top-down approaches are able to scale well, but their reliance on aggregate values and historical data makes them rigid, with limited capacity to view emergent behaviours or model scenarios that have not occurred historically. The resolution of these methods is also low, with modellers only able to observe aggregate behaviours, making it difficult to account for and measure the heterogeneity of households and their dwellings. The use of bottom-up modelling approaches rectified many of these problems, allowing the simulation of individual components from which aggregate values can be derived. However, such methods do not scale well due to the requirement to perform optimisation-based sub-problems which are too computationally costly to perform across an entire building stock. Some mixed approaches have been used in the literature, dividing the stock into a limited number of archetypes. These archetypes are used as representative dwelling-household combinations which are then upscaled to estimate the rest of the stock. These methods benefit from many of the bottom-up modelling approaches while remaining technically feasible, however, they do so at the cost of reducing the granularity of the model significantly as each household must conform to one of a limited number of selected archetypes, limiting the degree of heterogeneity. As the number of archetypes is derived from the combination of attributes of interest, modellers are required to keep the number of household attributes small to limit archetype numbers. The final hybrid approach found in the literature is the use of predictive machine learning models to replace the resource intensive sub-components of bottom-up modelling, expanding its feasibility. While this has been explored thoroughly when estimating energy demand, with a process tool referred to as a Surrogate Energy Performance Model (SEPM), very little research had been conducted into predictive models to replace the optimisation sub-component, a process referred to in this work as Surrogate Optimisation (SO). The only example found in the literature had not been applied to bottom-up retrofit adoption modelling, and was limited in the capacity for WHRSs-adoption modelling by the prediction of fixed-size and opaque Pareto fronts.

In order to tackle the research gap for technically feasible retrofit decision models with high resolution heterogeneous agents, the following research aim was constructed:

To expand the simulation possibility frontier using the Computer Science toolkit, to allow more descriptive bottom-up large-scale retrofit decision modelling without impacting the technical feasibility of the model.

Where the simulation possibility frontier represents the trade-offs that modellers must make between the technical feasibility of a simulation and the level of detail that simulation includes. In order to achieve this aim, several statements were laid out in the introduction which corresponded to significant contributions that this work intended to make to the literature. These contribution statements are repeated below, with reference to the chapter where the work to achieve them was primarily presented.

Contribution 1 - To the best knowledge of the author, this thesis is the first to systematically investigate the integration of Surrogate Optimisation into an Agent-Based Model to analyse energy retrofit adoption in urban housing stock. The integration of the SO technique into an ABM allows for increasingly rational, self-interested agents to be simulated at a scale that would otherwise be infeasible by allowing computationally cheap optimisations. The investigation in this thesis considers the implementation challenges, performance, and drawbacks of this method of ABM analysis. The work pertaining to this contribution was laid out in Chapter 5.

Contribution 2 - To the best knowledge of the author, this thesis is the first to systematically investigate the extension of the principle of multi-objective Surrogate Optimisation for the analysis of domestic urban energy retrofit potential that includes a measure of households' Willingness to Pay for carbon mitigation. This has allowed for the conception of the ABMs created for Contribution 1 to relax the assumption of self-interest by considering environmentally conscious agents. The work pertaining to this contribution was laid out in Chapter 6.

Contribution 3 - To the best knowledge of the author, this thesis is the first to combine a data-driven retrofit trigger model with a Surrogate Optimisation method. The decision trigger model takes survey data from participants to determine when retrofit adoptions are likely to be considered, as modelled in Contribution 2. This allows for models which conceptualise and include heterogeneous decision factors while maintaining intelligent and preference driven retrofit evaluations. The work pertaining to this contribution was laid out in Chapter 7.

While the contributions made in this work expand the feasibility and level of detail possible for urban scale WHRS-adoption modelling, it also opens up interesting avenues for possible future work. Much of this work can be considered in terms of the limitations or assumptions which could be lifted with extended work. This includes pushing out the scope and level of detail contained at the simulation, optimisation, or ABM stages of modelling, including features such as dynamic occupancy models, additional agent objectives, or extended HCV modelling. Additionally, future work could look to externally validate the retrofit decision model against real-world decisions, either by seeking to compile a data set of decisions made in the field or by using extended decision surveys which capture the impact of different factors on hypothetical decisions.

Appendix A

Additional Data

A.1 Building Data Set

Table A.1: Description of initial attributes of building data set.

Field Name	Description	Example	Notes
HSP	Heating Setpoint (Degrees Celsius)	21	Assigned Normal Distribution N(20.5, 2.5)
CurNCH	Nottingham City Homes	True	Boolean
RIR	Room in Roof	FALSE	True for $\tilde{1}1\%$ of Dwellings
hhType	Household Type	Lone Parent with Dependent Child	
TOID	Unique Topographic Identifier (Ordinance Survey)	osgb1000005074832	Primary Key in the Data Set
xStoreys	Number of Stories	2	
BHA_relhma	Relative Ridge Heights (Meters)	7.4	Ground to Peak of Roof
BHA_relh2	Relative Eave Heights (Meters)	5.7	Ground to Bottom of Roof
ECS	Heating Type (from ECS)	Boiler	
WType	Wall Type	Cavity	Cavity, Solid, or Other

Occupants	Estimated Number of Occupants at the Property	2	Mean: 2.3 Std. Deviation: 1.3
BldgUse	Building Usage	Residential	Residential, Vacant, or Mixed
TotalArea	Total Floor Area (Meters Squared)	78.31	Defined as footprint∗nstories
WallMat	Base Material Used in Wall (if known)	Brick	
WIns	Window Frame Insulation	TRUE	Boolean
EWR_form	Built Form of Building	Terraced	Detached, Semi-detached, Mid-terraced, end-terraced
geometry	Geometric Definition of Building Footprint Using Coordinates	Shapely Polygon	Shapely Python Geometry Object
WallGML	Reference to the GML Definition of Wall Construction	id_wall_5	Existing Construction (Based on EHS)
RoofGML	Reference to the GML Definition of Roof Construction	id_roof_1	Existing Construction (Based on EHS)
GlazGML	Reference to the GML Definition of Glazing type	id_glazing_4	Existing Construction (Based on EHS)
FrontGR	Glazing Ratio of Front of Building	20	Expressed as %
BackGR	Glazing Ratio of Back of Building	10	Expressed as %
SideGR	Glazing Ratio of Side of Building	15	Expressed as %
crAGE_DESC	Description of Age Classifier	1914-1945	9 in Total Including Non-residential Label
crAGE	Integer Representation of Age Classifier	1	1-9
LoftIns	Amount of Loft Insulation Present	100-150mm	4 Categories Ranging from 0-150mm+
crTYPE_DES	Building Archetype	Bungalow (Detached)	
crTYPE	Integer Representation of Building Archetype	2	13 Unique Types
InfilRate	Estimated Infiltration Rate	0.52	Air Changes per Hour
LSOA	Lower Layer Super Output Area (ONS)	Nottingham 021A	

MSOA	Middle Layer Super Output Area	Nottingham 021	
WardName	Administrative Ward the Property Belongs to	Clifton North	
oacode	Census Output Area (OA) Code	E00070449	

A.2 Optimisation Parameters

A.2.1 Cost Model Data

Type	Material	Thickness(mm)	Material Cost	Labour Cost	Notes
EWI	XPS	30	£14.50	£100.00	Per square meter
EWI	XPS	35	£15.45	£100.00	Per square meter
EWI	XPS	40	£16.41	£100.00	Per square meter
EWI	XPS	45	£17.36	£100.00	Per square meter
EWI	XPS	50	£18.32	£100.00	Per square meter
EWI	XPS	55	£19.28	£100.00	Per square meter
EWI	XPS	60	£20.23	£100.00	Per square meter
EWI	XPS	65	£21.19	£100.00	Per square meter
EWI	XPS	70	£22.14	£100.00	Per square meter
EWI	XPS	75	£23.10	£100.00	Per square meter
EWI	XPS	80	£24.06	£100.00	Per square meter
EWI	XPS	85	£25.01	£100.00	Per square meter
EWI	XPS	90	£25.97	£100.00	Per square meter
EWI	XPS	95	£26.92	£100.00	Per square meter
EWI	XPS	100	£27.88	£100.00	Per square meter
EWI	XPS	105	£28.84	£100.00	Per square meter
EWI	XPS	110	£29.79	£100.00	Per square meter
EWI	XPS	115	£30.75	£100.00	Per square meter
EWI	XPS	120	£31.70	£100.00	Per square meter
EWI	XPS	125	£32.66	£100.00	Per square meter
EWI	XPS	130	£33.62	£100.00	Per square meter
EWI	XPS	135	£34.57	£100.00	Per square meter
EWI	XPS	140	£35.53	£100.00	Per square meter
EWI	XPS	145	£36.48	£100.00	Per square meter
EWI	XPS	150	£37.44	£100.00	Per square meter

Type	Material	Thickness(mm)	Material Cost	Labour Cost	Notes
EWI	EPS	30	£11.67	£100.00	Per square meter
EWI	EPS	35	£12.27	£100.00	Per square meter
EWI	EPS	40	£12.86	£100.00	Per square meter
EWI	EPS	45	£13.46	£100.00	Per square meter
EWI	EPS	50	£14.05	£100.00	Per square meter
EWI	EPS	55	£14.65	£100.00	Per square meter
EWI	EPS	60	£15.24	£100.00	Per square meter
EWI	EPS	65	£15.84	£100.00	Per square meter
EWI	EPS	70	£16.43	£100.00	Per square meter
EWI	EPS	75	£17.03	£100.00	Per square meter
EWI	EPS	80	£17.62	£100.00	Per square meter
EWI	EPS	85	£18.22	£100.00	Per square meter
EWI	EPS	90	£18.81	£100.00	Per square meter
EWI	EPS	95	£19.41	£100.00	Per square meter
EWI	EPS	100	£20.00	£100.00	Per square meter
EWI	EPS	105	£20.60	£100.00	Per square meter
EWI	EPS	110	£21.19	£100.00	Per square meter
EWI	EPS	115	£21.79	£100.00	Per square meter
EWI	EPS	120	£22.38	£100.00	Per square meter
EWI	EPS	125	£22.98	£100.00	Per square meter
EWI	EPS	130	£23.57	£100.00	Per square meter
EWI	EPS	135	£24.17	£100.00	Per square meter
EWI	EPS	140	£24.76	£100.00	Per square meter
EWI	EPS	145	£25.36	£100.00	Per square meter
EWI	EPS	150	£25.100	£100.00	Per square meter
EWI	PIR	30	£12.99	£100.00	Per square meter
EWI	PIR	35	£13.43	£100.00	Per square meter
EWI	PIR	40	£13.87	£100.00	Per square meter
EWI	PIR	45	£14.31	£100.00	Per square meter
EWI	PIR	50	£14.76	£100.00	Per square meter
EWI	PIR	55	£15.20	£100.00	Per square meter
EWI	PIR	60	£15.64	£100.00	Per square meter
EWI	PIR	65	£16.08	£100.00	Per square meter
EWI	PIR	70	£16.53	£100.00	Per square meter
EWI	PIR	75	£16.97	£100.00	Per square meter
EWI	PIR	80	£17.41	£100.00	Per square meter
EWI	PIR	85	£17.86	£100.00	Per square meter

Type	Material	Thickness(mm)	Material Cost	Labour Cost	Notes
EWI	PIR	90	£18.30	£100.00	Per square meter
EWI	PIR	95	£18.74	£100.00	Per square meter
EWI	PIR	100	£19.18	£100.00	Per square meter
EWI	PIR	105	£19.63	£100.00	Per square meter
EWI	PIR	110	£20.07	£100.00	Per square meter
EWI	PIR	115	£20.51	£100.00	Per square meter
EWI	PIR	120	£20.100	£100.00	Per square meter
EWI	PIR	125	£21.40	£100.00	Per square meter
EWI	PIR	130	£21.84	£100.00	Per square meter
EWI	PIR	135	£22.28	£100.00	Per square meter
EWI	PIR	140	£22.73	£100.00	Per square meter
EWI	PIR	145	£23.17	£100.00	Per square meter
EWI	PIR	150	£23.61	£100.00	Per square meter
Roof	Mineral Wool	50	£3.26	£100.00	Per square meter
Roof	Mineral Wool	75	£4.86	£100.00	Per square meter
Roof	Mineral Wool	100	£6.47	£100.00	Per square meter
Roof	Mineral Wool	125	£8.07	£100.00	Per square meter
Roof	Mineral Wool	150	£9.67	£100.00	Per square meter
Roof	Mineral Wool	175	£11.28	£100.00	Per square meter
Roof	Mineral Wool	200	£12.88	£100.00	Per square meter
Roof	Mineral Wool	225	£14.48	£100.00	Per square meter
Roof	Mineral Wool	250	£16.09	£100.00	Per square meter
Roof	Mineral Wool	275	£17.69	£100.00	Per square meter
Roof	Mineral Wool	300	£19.29	£100.00	Per square meter
Roof	Mineral Wool	325	£20.89	£100.00	Per square meter
Roof	Mineral Wool	350	£22.50	£100.00	Per square meter
Roof	Mineral Wool	375	£24.10	£100.00	Per square meter
Roof	Mineral Wool	400	£25.70	£100.00	Per square meter
IWI	XPS	30	£19.26	£85.00	Per square meter
IWI	XPS	35	£20.21	£85.00	Per square meter
IWI	XPS	40	£21.17	£85.00	Per square meter
IWI	XPS	45	£22.12	£85.00	Per square meter
IWI	XPS	50	£23.08	£85.00	Per square meter
IWI	XPS	55	£24.04	£85.00	Per square meter
IWI	XPS	60	£24.99	£85.00	Per square meter
IWI	XPS	65	£25.100	£85.00	Per square meter
IWI	XPS	70	£26.90	£85.00	Per square meter

Type	Material	Thickness(mm)	Material Cost	Labour Cost	Notes
IWI	XPS	75	£27.86	£85.00	Per square meter
IWI	XPS	80	£28.82	£85.00	Per square meter
IWI	XPS	85	£29.77	£85.00	Per square meter
IWI	XPS	90	£30.73	£85.00	Per square meter
IWI	XPS	95	£31.68	£85.00	Per square meter
IWI	XPS	100	£32.64	£85.00	Per square meter
IWI	XPS	105	£33.60	£85.00	Per square meter
IWI	XPS	110	£34.55	£85.00	Per square meter
IWI	XPS	115	£35.51	£85.00	Per square meter
IWI	XPS	120	£36.46	£85.00	Per square meter
IWI	XPS	125	£37.42	£85.00	Per square meter
IWI	XPS	130	£38.38	£85.00	Per square meter
IWI	XPS	135	£39.33	£85.00	Per square meter
IWI	XPS	140	£40.29	£85.00	Per square meter
IWI	XPS	145	£41.24	£85.00	Per square meter
IWI	XPS	150	£42.20	£85.00	Per square meter
IWI	EPS	30	£16.43	£85.00	Per square meter
IWI	EPS	35	£17.03	£85.00	Per square meter
IWI	EPS	40	£17.62	£85.00	Per square meter
IWI	EPS	45	£18.22	£85.00	Per square meter
IWI	EPS	50	£18.81	£85.00	Per square meter
IWI	EPS	55	£19.41	£85.00	Per square meter
IWI	EPS	60	£20.00	£85.00	Per square meter
IWI	EPS	65	£20.60	£85.00	Per square meter
IWI	EPS	70	£21.19	£85.00	Per square meter
IWI	EPS	75	£21.79	£85.00	Per square meter
IWI	EPS	80	£22.38	£85.00	Per square meter
IWI	EPS	85	£22.98	£85.00	Per square meter
IWI	EPS	90	£23.57	£85.00	Per square meter
IWI	EPS	95	£24.17	£85.00	Per square meter
IWI	EPS	100	£24.76	£85.00	Per square meter
IWI	EPS	105	£25.36	£85.00	Per square meter
IWI	EPS	110	£25.10	£85.00	Per square meter
IWI	EPS	115	£26.55	£85.00	Per square meter
IWI	EPS	120	£27.14	£85.00	Per square meter
IWI	EPS	125	£27.74	£85.00	Per square meter
IWI	EPS	130	£28.33	£85.00	Per square meter

Type	Material	Thickness(mm)	Material Cost	Labour Cost	Notes
IWI	EPS	135	£28.93	£85.00	Per square meter
IWI	EPS	140	£29.52	£85.00	Per square meter
IWI	EPS	145	£30.12	£85.00	Per square meter
IWI	EPS	150	£30.71	£85.00	Per square meter
IWI	PIR	30	£17.75	£85.00	Per square meter
IWI	PIR	35	£18.19	£85.00	Per square meter
IWI	PIR	40	£18.63	£85.00	Per square meter
IWI	PIR	45	£19.07	£85.00	Per square meter
IWI	PIR	50	£19.52	£85.00	Per square meter
IWI	PIR	55	£19.96	£85.00	Per square meter
IWI	PIR	60	£20.40	£85.00	Per square meter
IWI	PIR	65	£20.84	£85.00	Per square meter
IWI	PIR	70	£21.29	£85.00	Per square meter
IWI	PIR	75	£21.73	£85.00	Per square meter
IWI	PIR	80	£22.17	£85.00	Per square meter
IWI	PIR	85	£22.62	£85.00	Per square meter
IWI	PIR	90	£23.06	£85.00	Per square meter
IWI	PIR	95	£23.50	£85.00	Per square meter
IWI	PIR	100	£23.94	£85.00	Per square meter
IWI	PIR	105	£24.39	£85.00	Per square meter
IWI	PIR	110	£24.83	£85.00	Per square meter
IWI	PIR	115	£25.27	£85.00	Per square meter
IWI	PIR	120	£25.71	£85.00	Per square meter
IWI	PIR	125	£26.16	£85.00	Per square meter
IWI	PIR	130	£26.60	£85.00	Per square meter
IWI	PIR	135	£27.04	£85.00	Per square meter
IWI	PIR	140	£27.49	£85.00	Per square meter
IWI	PIR	145	£27.93	£85.00	Per square meter
IWI	PIR	150	£28.37	£85.00	Per square meter
Heating	Electric Heating	N/A	£200.00	£-	Per Heating Zone
Heating	Gas Heating	N/A	£1,000.00	£2,000.00	Per Installation
Glazing	Single Glazing	N/A	£-	£-	Excluded (Retrofit)
Glazing	Double Glazing	N/A	£325.00	£100.00	Per Window
Glazing	Triple Glazing	N/A	£425.00	£100.00	Per Window

A.2.2 Genetic Algorithm Tuning

Table A.3: GA tuning results for NPV-maximising objective.

Mean NPV	NPV Variance	Opt Time (s)	Mutation Method	Recombination Method	Replacement Method	Mutation Rate	Pop Size
3948.99	52098	1.92	RANDOM	2 POINT	PURE ELITE	0.10	32
3948.96	64793	1.93	RANDOM	UNIFORM	PURE ELITE	0.10	32
3948.95	89811	1.90	RANDOM	2 POINT	PURE ELITE	0.05	32
3948.95	88791	1.90	RANDOM	UNIFORM	PURE ELITE	0.05	32
3948.95	127511	1.95	UNIFORM	UNIFORM	PURE ELITE	0.10	32
3948.89	141529	1.95	UNIFORM	2 POINT	PURE ELITE	0.05	32
3948.86	95667	2.15	RANDOM	2 POINT	SOFT ELITE	0.05	32
3948.78	123401	2.14	UNIFORM	2 POINT	SOFT ELITE	0.10	32
3948.77	130700	1.94	UNIFORM	UNIFORM	PURE ELITE	0.05	32
3948.76	126096	1.97	UNIFORM	2 POINT	PURE ELITE	0.10	32
3948.50	68144	2.14	RANDOM	2 POINT	SOFT ELITE	0.10	32
3948.25	77376	2.18	RANDOM	UNIFORM	SOFT ELITE	0.10	32
3948.18	173949	2.03	UNIFORM	2 POINT	SOFT ELITE	0.05	32
3948.06	161521	1.15	RANDOM	2 POINT	PURE ELITE	0.10	16
3947.90	190432	1.18	RANDOM	UNIFORM	PURE ELITE	0.10	16
3947.86	87762	2.14	RANDOM	UNIFORM	SOFT ELITE	0.05	32
3947.52	122772	2.13	UNIFORM	UNIFORM	SOFT ELITE	0.10	32
3947.36	520873	1.17	UNIFORM	UNIFORM	PURE ELITE	0.10	16
3947.33	273904	1.13	RANDOM	2 POINT	PURE ELITE	0.05	16
3947.07	209531	1.34	RANDOM	2 POINT	SOFT ELITE	0.10	16
3946.91	533635	1.22	UNIFORM	2 POINT	PURE ELITE	0.10	16
3946.85	381169	1.14	RANDOM	UNIFORM	PURE ELITE	0.05	16
3946.51	650612	1.14	UNIFORM	2 POINT	PURE ELITE	0.05	16
3946.47	242629	1.37	RANDOM	UNIFORM	SOFT ELITE	0.10	16
3946.39	146410	2.06	UNIFORM	UNIFORM	SOFT ELITE	0.05	32
3946.25	322628	1.28	RANDOM	UNIFORM	SOFT ELITE	0.05	16
3945.37	377354	1.25	RANDOM	2 POINT	SOFT ELITE	0.05	16
3943.86	535633	1.29	UNIFORM	2 POINT	SOFT ELITE	0.10	16
3943.68	549545	1.15	UNIFORM	UNIFORM	PURE ELITE	0.05	16
3941.04	442749	1.27	UNIFORM	UNIFORM	SOFT ELITE	0.10	16
3940.10	850229	0.77	RANDOM	2 POINT	PURE ELITE	0.10	8
3939.98	606333	1.20	UNIFORM	2 POINT	SOFT ELITE	0.05	16
3939.92	721621	0.76	RANDOM	UNIFORM	PURE ELITE	0.10	8
3938.61	608375	1.23	UNIFORM	UNIFORM	SOFT ELITE	0.05	16
3937.11	1212386	0.83	RANDOM	2 POINT	SOFT ELITE	0.10	8
3934.92	1246482	0.71	RANDOM	UNIFORM	PURE ELITE	0.05	8
3933.30	834979	0.82	RANDOM	UNIFORM	SOFT ELITE	0.10	8
3932.89	1200359	0.73	RANDOM	2 POINT	PURE ELITE	0.05	8
3930.30	1535799	0.71	UNIFORM	UNIFORM	PURE ELITE	0.10	8
3927.97	1628791	0.74	UNIFORM	2 POINT	PURE ELITE	0.10	8
3926.51	1368350	0.77	RANDOM	2 POINT	SOFT ELITE	0.05	8
3925.56	1083287	0.76	RANDOM	UNIFORM	SOFT ELITE	0.05	8
3917.79	1698322	0.76	UNIFORM	UNIFORM	SOFT ELITE	0.10	8
3917.41	1768710	0.80	UNIFORM	2 POINT	SOFT ELITE	0.10	8
3917.20	1731973	0.73	UNIFORM	2 POINT	PURE ELITE	0.05	8
3912.30	1829741	0.69	UNIFORM	UNIFORM	PURE ELITE	0.05	8

A.3 Additional Surrogate Optimiser Results

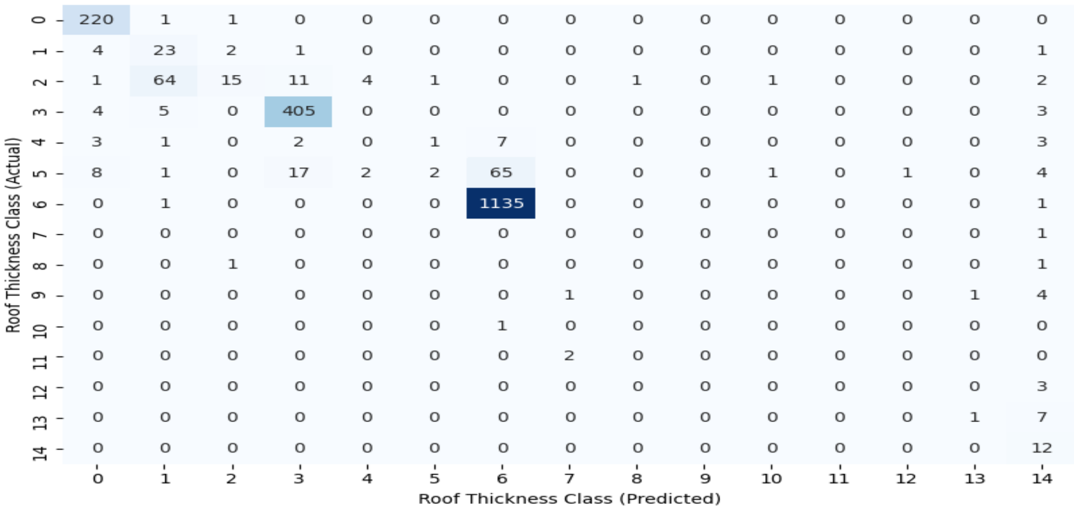

Figure A.1: Confusion matrix of single-objective Surrogate Optimiser roof insulation thickness results coerced back into thickness classes.

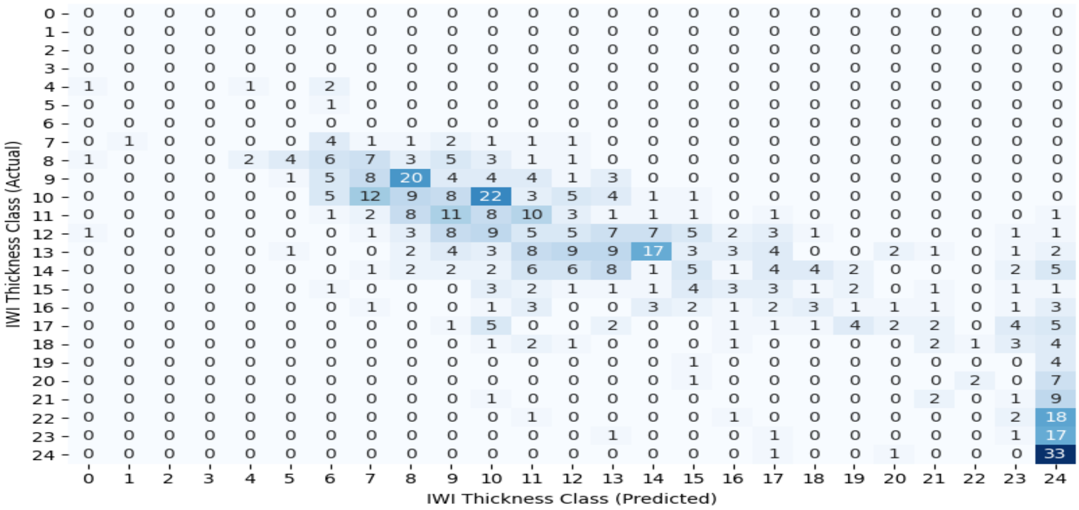

Figure A.2: Confusion matrix of single-objective Surrogate Optimiser IWI insulation thickness results coerced back into thickness classes.

References

[1] Jason Palmer, Mary Livingstone, and Angela Adams. What Does It Cost To Retrofit Homes? Technical report, Department for Business, Energy and Industrial Strategy, 2017. URL: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/656866/BEIS_Update_of_Domestic_Cost_Assumptions_031017.pdf.

[2] Naomi Oreskes. The scientific consensus on climate change. *Science*, 306(5702):1686, 12 2004. doi:10.1126/science.1103618.

[3] Peter T. Doran and Maggie Kendall Zimmerman. Examining the scientific consensus on climate change. *Eos*, 90(3):22–23, 1 2011. doi:10.1029/2009eo030002.

[4] Mark Lynas, Benjamin Z. Houlton, and Simon Perry. Greater than 99% consensus on human caused climate change in the peer-reviewed scientific literature. *Environmental Research Letters*, 16(11):114005, 10 2021. doi:10.1088/1748-9326/AC2966.

[5] The UK Government. Climate Change Act 2008, 2008. URL: https://www.legislation.gov.uk/ukpga/2008/27/section/2.

[6] Waters L. Energy Consumption in the UK (ECUK) 1970 to 2019. Technical report, Office for National Statistics, 2020. URL: https://www.gov.uk/government/collections/digest-of-uk-energy-statistics-dukes.

[7] Office For National Statistics. Provisional UK greenhouse gas emissions national statistics, 2021. URL: https://data.gov.uk/dataset/9a1e58e5-d1b6-457d-a414-335ca546d52c/provisional-uk-greenhouse-gas-emissions-national-statistics.

[8] Michael J. North. A theoretical formalism for analyzing agent-based models. *Complex Adaptive Systems Modeling*, 2(1):1–34, 12 2014. doi:10.1186/2194-3206-2-3.

[9] Deborah Wilson and Joel Swisher. Exploring the gap. Top-down versus bottom-up analyses of the cost of mitigating global warming. *Energy Policy*, 21(3):249–263, 3 1993. doi:10.1016/0301-4215(93)90247-D.

[10] Peer Olaf Siebers, Zhi En Lim, Grazziela P. Figueredo, and James Hey. An innovative approach to multi-method integrated assessment modelling of global climate change. *JASSS*, 23(1), 1 2020. doi:10.18564/jasss.4209.

[11] Eric Bonabeau. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the United States of America*, 99(SUPPL. 3):7280–7287, 5 2002. doi:10.1073/pnas.082080899.

[12] Li An. Modeling human decisions in coupled human and natural systems: Review of agent-based models. *Ecological Modelling*, 229:25–36, 2012. doi:10.1016/j.ecolmodel.2011.07.010.

[13] Muaz Niazi and Amir Hussain. Agent-based computing from multi-agent systems to agent-based models: A visual survey. *Scientometrics*, 89(2):479–499, 2011. doi:10.1007/s11192-011-0468-9.

[14] Stuart J Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*, volume 1. Prentice Hall, Hoboken, NJ, 1995. URL: https://dl.acm.org/doi/book/10.5555/773294.

[15] Andrew T. Crooks and Alison J. Heppenstall. Introduction to agent-based modelling. In *Agent-Based Models of Geographical Systems*, pages 85–105. Springer Netherlands, 1 2012. URL: https://link.springer.com/chapter/10.1007/978-90-481-8927-4_5, doi:10.1007/978-90-481-8927-4{_}5/COVER/.

[16] Jang Won Bae, Euihyun Paik, Kiho Kim, Karandeep Singh, and Mazhar Sajjad. Combining Microsimulation and Agent-based Model for Micro-level Population Dynamics. *Procedia Computer Science*, 80:507–517, 1 2016. doi:10.1016/J.PROCS.2016.05.331.

[17] Mirjana Pejic Bach, Emil Tustanovski, Andrew W.H. Ip, Kai Leung Yung, and Vasja Roblek. System dynamics models for the simulation of sustainable urban development: A review and analysis and the stakeholder perspective. *Kybernetes*, 49(2):460–504, 1 2020. URL: https://www.emerald.com/insight/content/doi/10.1108/K-04-2018-0210/full/pdf.

[18] Danilo Bzdok, Naomi Altman, and Martin Krzywinski. Statistics versus machine learning. *Nature Methods*, 15:233–234, 2018. doi:10.1038/nmeth.4642.

[19] Ethem Alpaydin. Introduction. In *Introduction to Machine Learning*, pages 1–19. MIT Press, 2010. URL: `https://ieeexplore.ieee.org/book/6267367`.

[20] Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien. *Semi-Supervised Learning*. MIT Press, Cambridge, MA, 2006. URL: `https://ieeexplore.ieee.org/book/6267236`.

[21] Johannes Dahlke, Kristina Bogner, Matthias Mueller, Thomas Berger, Andreas Pyka, and Bernd Ebersberger. Is the Juice Worth the Squeeze? Machine Learning (ML) In and For Agent-Based Modelling (ABM). *arXiv*, 3 2020. URL: `http://arxiv.org/abs/2003.11985`.

[22] Christopher J.C. Burges. A Tutorial on Support Vector Machines for Pattern Recognition. *Data Mining and Knowledge Discovery 1998 2:2*, 2(2):121–167, 1998.

[23] N. S. Altman. An Introduction to Kernel and Nearest-Neighbor Nonparametric Regression. *The American Statistician*, 46(3):175, 8 1992. `doi:10.2307/2685209`.

[24] Blake Shaw and Tony Jebara. Structure preserving embedding. In *ACM International Conference Proceeding Series*, volume 382, 2009. `doi:10.1145/1553374.1553494`.

[25] Jürgen Schmidhuber. Deep Learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015. `doi:10.1016/j.neunet.2014.09.003`.

[26] M D Zeiler, M Ranzato, R Monga, M Mao, K Yang, Q V Le, P Nguyen, A Senior, V Vanhoucke, J Dean, and G E Hinton. On Rectified Linear Units For Speech Processing. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 3517–3521, 2013. `doi:10.1109/ICASSP.2013.6638312`.

[27] J. MacQueen. Some methods for classification and analysis of multivariate observations. *Berkeley Symposium on Mathematical Statistics and Probability*, 1:281–297, 1 1967.

[28] J. C. Dunn. A Fuzzy Relative of the ISODATA Process and Its Use in Detecting Compact Well-Separated Clusters. *Journal of Cybernetics*, 3(3):32–57, 1 1973. `doi:10.1080/01969727308546046`.

[29] James C. Bezdek, Robert Ehrlich, and William Full. FCM: The fuzzy c-means clustering algorithm. *Computers & Geosciences*, 10(2-3):191–203, 1 1984. `doi:10.1016/0098-3004(84)90020-7`.

[30] Kenneth Sörensen and Fred W. Glover. Metaheuristics. In *Encyclopedia of Operations Research and Management Science*, pages 960–970. Springer, Boston,

MA, 2013. URL: https://link.springer.com/referenceworkentry/10.1007/978-1-4419-1153-7_1167.

[31] Kenneth Sörensen, Marc Sevaux, and Fred Glover. A History of Metaheuristics. In *Handbook of Heuristics*. Springer International Publishing, Cham, 2018. URL: https://link.springer.com/referenceworkentry/10.1007/978-3-319-07153-4_4-1.

[32] E.H.L Aarts and J.H.M. Korst. *Simulated Annealing and Boltzmann Machines: A Stochastic Approach to Combinatorial Optimization and Neural Computing*. John Wiley & Sons, Chichester, 1989.

[33] Alex Kwaku Peprah, Simon Kojo Appiah, Samuel Kwame Amponsah, Alex Kwaku Peprah, Simon Kojo Appiah, and Samuel Kwame Amponsah. An Optimal Cooling Schedule Using a Simulated Annealing Based Approach. *Applied Mathematics*, 8(8):1195–1210, 8 2017. doi:10.4236/AM.2017.88090.

[34] David E Goldberg. *Genetic Algorithms in Search, Optimization & Machine Learning*. Addison-Wesley Publishing Company, INC., Reading, Massachusetts, 1989.

[35] A. Messac, A. Ismail-Yahaya, and C.A. Mattson. The normalized normal constraint method for generating the Pareto frontier. *Structural and Multidisciplinary Optimization*, 25(2):86–98, 6 2003. doi:10.1007/S00158-002-0276-1.

[36] George Mavrotas. Effective implementation of the ϵ-constraint method in Multi-Objective Mathematical Programming problems. *Applied Mathematics and Computation*, 213(2):455–465, 7 2009. doi:10.1016/J.AMC.2009.03.037.

[37] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. A Fast and Elitist Multi-Objective Genetic Algorithm: NSGA-II. *IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION*, 6(2), 2000. URL: https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.17.7771.

[38] Hisao Ishibuchi, Naoya Akedo, and Yusuke Nojima. Behavior of Multiobjective Evolutionary Algorithms on Many-Objective Knapsack Problems. *IEEE Transactions on Evolutionary Computation*, 19(2):264–283, 4 2015. doi:10.1109/TEVC.2014.2315442.

[39] Drury B Crawley, Curtis O Pedersen, Linda K Lawrie, and Frederick C Winkelmann. EnergyPlus : Energy Simulation Program. *ASHRAE Journal*, 1(April):49–56, 2000. URL: https://www.researchgate.net/profile/Drury-Crawley/publication/230606369_EnergyPlus_Energy_Simulation_Program/links/546a40d80cf2f5eb18077919/EnergyPlus-Energy-Simulation-Program.

[40] EnergyPlus. EnergyPlus — EnergyPlus, 2021. URL: `https://energyplus.net/`.

[41] Darren Robinson, Frédéric Haldi, Jérôme Henri Kämpf, P Leroux, D Perez, a Rasheed, and U Wilke. CITYSIM: Comprehensive Micro-Simulation Of Resource Flows For Sustainable Urban Planning. *International IBPSA Conference*, pages 1083–1090, 2009.

[42] U.S. Department of Energy. EnergyPlus Documentation: Getting Started with EnergyPlus Basic Concepts. Technical report, U.S. Department of Energy, 2021. URL: `https://energyplus.net/assets/nrel_custom/pdfs/pdfs_v9.6.0/GettingStarted.pdf`.

[43] Glenn W. Harrison. Expected Utility Theory and the Experimentalists. In *Experimental Economics*, pages 43–73. Physica-Verlag HD, Heidelberg, 1994. URL: `https://link.springer.com/chapter/10.1007/978-3-642-51179-0_4`.

[44] Glenn W. Harrison and E. Elisabet Rutström. Risk Aversion in the Laboratory. *Research in Experimental Economics*, 12:41–196, 2008. `doi:10.1016/S0193-2306(08)00003-3`.

[45] A Tversky and D Kahneman. The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458, 1 1981. `doi:10.1126/SCIENCE.7455683`.

[46] Anton Kühberger. The Influence of Framing on Risky Decisions: A Meta-analysis. *Organizational Behavior and Human Decision Processes*, 75(1):23–55, 7 1998. `doi:10.1006/OBHD.1998.2781`.

[47] Sascha Hokamp and Michael Pickhardt. Income Tax Evasion in a Society of Heterogeneous Agents-Evidence from an Agent-based Model. *International Economic Journal*, 24(4):541–553, 2010. `doi:10.1080/10168737.2010.525994`.

[48] Chaohui Zheng, Yi Liu, Bettina Bluemling, Jining Chen, and Arthur P.J. Mol. Modeling the environmental behavior and performance of livestock farmers in China: An ABM approach. *Agricultural Systems*, 122:60–72, 11 2013. `doi:10.1016/J.AGSY.2013.08.005`.

[49] Kim M. Bloomquist. A Comparison of Agent-Based Models of Income Tax Evasion:. *Social Science Computer Review*, 24(4):411–425, 8 2016. `doi:10.1177/0894439306287021`.

[50] Ted O'Donoghue and Jason Somerville. Modeling Risk Aversion in Economics. *Journal of Economic Perspectives*, 32(2):91–114, 3 2018. `doi:10.1257/JEP.32.2.91`.

[51] James Andreoni. Why free ride? *Journal of Public Economics*, 37(3):291–304, 1988. doi:10.1016/0047-2727(88)90043-6.

[52] James Andreoni. Impure Altruism and Donations to Public Goods : A Theory of Warm-Glow Giving. *The Economic Journal*, 100(401):464–477, 1990. doi:10.2307/2234133.

[53] Bruno S. Frey and Stephan Meier. Social Comparisons and Pro-social Behavior: Testing Conditional Cooperation in a Field Experiment. *American Economic Review*, 94(5):1717–1722, 12 2004. doi:10.1257/0002828043052187.

[54] Y. H. Farzin, K. J.M. Huisman, and P. M. Kort. Optimal timing of technology adoption. *Journal of Economic Dynamics and Control*, 22(5):779–799, 1998. doi:10.1016/s0165-1889(97)00097-3.

[55] D S Remer and A P Nieto. A compendium and comparison of 25 project evaluation techniques. Part 1: Net present value and rate of return methods. *Int. J. Production Economics 42*, 42:79–96, 1995. URL: https://www.sciencedirect.com/science/article/pii/0925527395001042.

[56] Ondřej Žižlavský. Net Present Value Approach: Method for Economic Assessment of Innovation Projects. *Procedia - Social and Behavioral Sciences*, 156(April):506–512, 2014. doi:10.1016/j.sbspro.2014.11.230.

[57] Ulrich Doraszelski. The net present value method versus the option value of waiting : A note on. *Journal of Economic Dynamics & Control*, 25(2001):1109–1115, 2009.

[58] Energy Industrial Strategy Department for Business. Standard Assessment Procedure, 2021. URL: https://www.gov.uk/guidance/standard-assessment-procedure.

[59] U.S. Department of Energy. Home Energy Score Scoring Methodology. Technical report, U.S. Department of Energy, 2021. URL: www.HomeEnergyScore.gov.

[60] J Henderson and J Hart. BREDEM 2012-A technical description of the BRE Domestic Energy Model. Technical report, BRE, 2013. URL: https://www.bre.co.uk/filelibrary/bredem/BREDEM-2012-specification.pdf.

[61] The Department of Energy and Climate Change. SAP 2012 The Governments Standard Assessment Procedure for Energy Rating of Dwellings 2012 edition. Technical report, BREgroup, 2014. URL: www.bre.co.uk/sap2012.

[62] Scott Kelly, Doug Crawford-Brown, and Michael G. Pollitt. Building performance evaluation and certification in the UK: Is SAP fit for purpose? *Renewable and*

Sustainable Energy Reviews, 16(9):6861–6878, 12 2012. doi:10.1016/J.RSER.2012.07.018.

[63] Slawomir Koziel, Qingsha S. Cheng, and John W. Bandler. Space mapping. *IEEE Microwave Magazine*, 9(6):105–122, 2008. doi:10.1109/MMM.2008.929554.

[64] Bryan Eisenhower, Zheng O'Neill, Satish Narayanan, Vladimir A. Fonoberov, and Igor Mezić. A methodology for meta-model based optimization in building energy models. *Energy and Buildings*, 47:292–301, 4 2012. doi:10.1016/j.enbuild.2011.12.001.

[65] Ralph Evins. A review of computational optimisation methods applied to sustainable building design. *Renewable and Sustainable Energy Reviews*, 22:230–245, 2013. doi:10.1016/j.rser.2013.02.004.

[66] Arfa Nawal Aijazi. *Machine Learning Paradigms for Building Energy Performance Simulations*. PhD thesis, Massachusetts Institute of Technology, 2017.

[67] Shreshth Nagpal, Caitlin Mueller, Arfa Aijazi, and Christoph F. Reinhart. A methodology for auto-calibrating urban building energy models using surrogate modeling techniques. *Journal of Building Performance Simulation*, 12(1):1–16, 1 2018. doi:10.1080/19401493.2018.1457722.

[68] Stavros Tseranidis, Nathan C. Brown, and Caitlin T. Mueller. Data-driven approximation algorithms for rapid performance evaluation and optimization of civil structures. *Automation in Construction*, 72:279–293, 12 2016. doi:10.1016/J.AUTCON.2016.02.002.

[69] Richard E. Edwards, Joshua New, Lynne E. Parker, Borui Cui, and Jin Dong. Constructing large scale surrogate models from big data and artificial intelligence. *Applied Energy*, 202:685–699, 9 2017. doi:10.1016/j.apenergy.2017.05.155.

[70] Wei Tian, Ruchi Choudhary, Godfried Augenbroe, and Sang Hoon Lee. Importance analysis and meta-model construction with correlated variables in evaluation of thermal performance of campus buildings. *Building and Environment*, 92:61–74, 2015. doi:10.1016/j.buildenv.2015.04.021.

[71] Paul Westermann and Ralph Evins. Using Bayesian deep learning approaches for uncertainty-aware building energy surrogate models. *Energy and AI*, 3:100039, 3 2021. doi:10.1016/J.EGYAI.2020.100039.

[72] P. Wate, M. Iglesias, V. Coors, and D. Robinson. Framework for emulation and uncertainty quantification of a stochastic building performance simulator. *Applied Energy*, 258, 1 2020. doi:10.1016/j.apenergy.2019.113759.

[73] Laurent Magnier and Fariborz Haghighat. Multiobjective optimization of building design using TRNSYS simulations, genetic algorithm, and Artificial Neural Network. *Building and Environment*, 45(3):739–746, 3 2010. doi:10.1016/J.BUILDENV.2009.08.016.

[74] Fabrizio Ascione, Nicola Bianco, Claudio De Stasio, Gerardo Maria Mauro, and Giuseppe Peter Vanoli. CASA, cost-optimal analysis by multi-objective optimisation and artificial neural networks: A new framework for the robust assessment of cost-optimal energy retrofit, feasible for any building. *Energy and Buildings*, 146:200–219, 7 2017. doi:10.1016/j.enbuild.2017.04.069.

[75] A. Prada, A. Gasparella, and P. Baggio. On the performance of meta-models in building design optimization. *Applied Energy*, 225:814–826, 9 2018. doi:10.1016/j.apenergy.2018.04.129.

[76] Christoph Waibel, Thomas Wortmann, Ralph Evins, and Jan Carmeliet. Building energy optimization: An extensive benchmark of global search algorithms. *Energy and Buildings*, 187:218–240, 3 2019. doi:10.1016/J.ENBUILD.2019.01.048.

[77] Seyed Amirhosain Sharif and Amin Hammad. Developing surrogate ANN for selecting near-optimal building energy renovation methods considering energy consumption, LCC and LCA. *Journal of Building Engineering*, 25:100790, 9 2019. doi:10.1016/j.jobe.2019.100790.

[78] Paul Westermann, Theodor Victor Christiaanse, Will Beckett, Paul Kovacs, and Ralph Evins. besos: Building and Energy Simulation, Optimization and Surrogate Modelling. *Journal of Open Source Software*, 6(60):2677, 2021. doi:10.21105/joss.02677.

[79] Paul Westermann and Ralph Evins. Surrogate modelling for sustainable building design – A review. *Energy and Buildings*, 198:170–186, 9 2019. doi:10.1016/J.ENBUILD.2019.05.057.

[80] Fabrizio Ascione, Nicola Bianco, Claudio De Stasio, Gerardo Maria Mauro, and Giuseppe Peter Vanoli. Artificial neural networks to predict energy performance and retrofit scenarios for any member of a building category: A novel approach. *Energy*, 118:999–1017, 1 2017. doi:10.1016/J.ENERGY.2016.10.126.

[81] James Hey, Peer-Olaf Siebers, Paul Nathanail, Ender Özcan, and Darren Robinson. Surrogate Optimisation of Housing Stock Retrofits using Deep Neural Networks. In *Proceedings of the 5th IBPSA-England Conference on Building Simulation and Optimization (Virtual)*, 2020.

[82] Emmanouil Thrampoulidis, Georgios Mavromatidis, Aurelien Lucchi, and Kristina Orehounig. A machine learning-based surrogate model to approximate optimal building retrofit solutions. *Applied Energy*, 281:116024, 1 2021. `doi:10.1016/j.apenergy.2020.116024`.

[83] Alberto Costa, Giacomo Nannicini, A Costa, Eth Zurich, and G Nannicini. RBFOpt: an open-source library for black-box optimization with costly function evaluations. *Mathematical Programming Computation*, 10:597–629, 2018. `doi:10.5281/zenodo.597767`.

[84] J. R. Caldwell, R. A. Watson, C. Thies, and J. D. Knowles. Deep Optimisation: Solving Combinatorial Optimisation Problems using Deep Neural Networks. *arXiv*, 11 2018. URL: `https://arxiv.org/abs/1811.00784v1`.

[85] Thomas Wortmann. Genetic evolution vs. function approximation: Benchmarking algorithms for architectural design optimization. *Journal of Computational Design and Engineering*, 6:414–428, 2018. `doi:10.1016/j.jcde.2018.09.001`.

[86] Lukas G. Swan and V. Ismet Ugursal. Modeling of end-use energy consumption in the residential sector: A review of modeling techniques. *Renewable and Sustainable Energy Reviews*, 13(8):1819–1835, 10 2009. `doi:10.1016/J.RSER.2008.09.033`.

[87] Lisa M.H. Hall and Alastair R. Buckley. A review of energy systems models in the UK: Prevalent usage and categorisation. *Applied Energy*, 169:607–628, 5 2016. `doi:10.1016/J.APENERGY.2016.02.044`.

[88] Gustavo Sousa, Benjamin M. Jones, Parham A. Mirzaei, and Darren Robinson. A review and critique of UK housing stock energy models, modelling approaches and data sources. *Energy and Buildings*, 151:66–80, 9 2017. `doi:10.1016/j.enbuild.2017.06.043`.

[89] M. Kavgic, A. Mavrogianni, D. Mumovic, A. Summerfield, Z. Stevanovic, and M. Djurovic-Petrovic. A review of bottom-up building stock models for energy consumption in the residential sector. *Building and Environment*, 45(7):1683–1697, 7 2010. `doi:10.1016/j.buildenv.2010.01.021`.

[90] J. Langevin, J. L. Reyna, S. Ebrahimigharehbaghi, N. Sandberg, P. Fennell, C. Nägeli, J. Laverge, M. Delghust, Mata, M. Van Hove, J. Webster, F. Federico, M. Jakob, and C. Camarasa. Developing a common approach for classifying building stock energy models. *Renewable and Sustainable Energy Reviews*, 133, 11 2020. `doi:10.1016/J.RSER.2020.110276`.

[91] Merih Aydinalp-Koksal and V. Ismet Ugursal. Comparison of neural network, conditional demand analysis, and engineering approaches for modeling end-use

energy consumption in the residential sector. *Applied Energy*, 85(4):271–296, 4 2008. doi:10.1016/J.APENERGY.2006.09.012.

[92] Timothy Lee and Runming Yao. Incorporating technology buying behaviour into UK-based long term domestic stock energy models to provide improved policy analysis. *Energy Policy*, 52:363–372, 1 2013. doi:10.1016/J.ENPOL.2012.09.048.

[93] M Hughes, J Palmer, and P Pope. A Guide to The Cambridge Housing Model. Technical report, Department of Energy Climate Change (DEC), 11 2013.

[94] Gustavo Sousa, Benjamin M. Jones, Parham A. Mirzaei, and Darren Robinson. An open-source simulation platform to support the formulation of housing stock decarbonisation strategies. *Energy and Buildings*, 172:459–477, 8 2018. doi:10.1016/J.ENBUILD.2018.05.015.

[95] Roger Amstalden, Michael Kost, Carsten Nathani, and Dieter Imboden. Economic potential of energy-efficient retrofitting in the Swiss residential building sector: The effects of policy instruments and energy price expectations. *Energy Policy*, 35(3):1819–1829, 2007. doi:10.1016/J.ENPOL.2006.05.018.

[96] A. P. Melo, D. Cóstola, R. Lamberts, and J. L.M. Hensen. Development of surrogate models using artificial neural network for building shell energy labelling. *Energy Policy*, 69:457–466, 2014. doi:10.1016/j.enpol.2014.02.001.

[97] Alessio Mastrucci, Olivier Baume, Francesca Stazi, and Ulrich Leopold. Estimating energy savings for the residential building stock of an entire city: A GIS-based statistical downscaling approach applied to Rotterdam. *Energy and Buildings*, 75:358–367, 2014. doi:10.1016/J.ENBUILD.2014.02.032.

[98] Alaia Sola, Cristina Corchero, Jaume Salom, and Manel Sanmarti. Simulation Tools to Build Urban-Scale Energy Models: A Review. *Energies*, 11(12):3269, 11 2018. URL: http://www.mdpi.com/1996-1073/11/12/3269, doi:10.3390/en11123269.

[99] J Huber and C Nytsch-Geusen. Development of modeling and simulation strategies for large-scale urbandistricts. In *Conference of International Building Performance Simulation Association*, pages 1753–1760, Sydney, Australia, 11 2011. URL: http://ibpsa.org/proceedings/BS2011/P_1569.pdf.

[100] Thomas H. Kolbe, Gerhard Gröger, and Lutz Plümer. CityGML: Interoperable Access to 3D City Models. *Geo-information for Disaster Management*, pages 883–899, 2005. URL: https://link.springer.com/chapter/10.1007/3-540-27468-5_63.

[101] Thomas H. Kolbe. Representing and Exchanging 3D City Models with CityGML. *Lecture Notes in Geoinformation and Cartography*, pages 15–31, 2009. URL: https://link.springer.com/chapter/10.1007/978-3-540-87395-2_2.

[102] Marcus Goetz. Towards generating highly detailed 3D CityGML models from OpenStreetMap. *International Journal of Geographical Information Science*, 27(5):845–865, 5 2013. doi:10.1080/13658816.2012.721552.

[103] Julian F. Rosser, Gavin Long, Sameh Zakhary, Doreen S. Boyd, Yong Mao, and Darren Robinson. Modelling urban housing stocks for building energy simulation using CityGML energyADE. *ISPRS International Journal of Geo-Information*, 8(4):163, 3 2019. URL: www.mdpi.com/journal/ijgi, doi:10.3390/ijgi8040163.

[104] Maikel Issermann, Fi John Chang, and Pu Yun Kow. Interactive urban building energy modelling with functional mockup interface of a local residential building stock. *Journal of Cleaner Production*, 289:125683, 3 2021. doi:10.1016/J.JCLEPRO.2020.125683.

[105] Oleksii Pasichnyi, Jörgen Wallin, and Olga Kordas. Data-driven building archetypes for urban building energy modelling. *Energy*, 181:360–377, 8 2019. doi:10.1016/J.ENERGY.2019.04.197.

[106] Carol C. Menassa and Brad Baer. A framework to assess the role of stakeholders in sustainable building retrofit decisions. *Sustainable Cities and Society*, 10:207–221, 2 2014. doi:10.1016/j.scs.2013.09.002.

[107] Taehoon Hong, Choongwan Koo, Hyunjoong Kim, and Hyo Seon Park. Decision support model for establishing the optimal energy retrofit strategy for existing multi-family housing complexes. *Energy Policy*, 66:157–169, 3 2014. doi:10.1016/J.ENPOL.2013.10.057.

[108] Mohammad Pazouki, Kamran Rezaie, and Ali Bozorgi-Amiri. A fuzzy robust multi-objective optimization model for building energy retrofit considering utility function: A university building case study. *Energy and Buildings*, 241:110933, 6 2021. doi:10.1016/J.ENBUILD.2021.110933.

[109] D. Duah and M. Syal. Intelligent decision support system for home energy retrofit adoption. *International Journal of Sustainable Built Environment*, 5(2):620–634, 12 2016. doi:10.1016/J.IJSBE.2016.05.003.

[110] D'Alpaos Chiara and Paolo Bragolusi. Buildings energy retrofit valuation approaches: state of the art and future perspectives. *Valori e Valutazioni*, (20), 2018.

[111] Pernilla Gluch and Henrikke Baumann. The life cycle costing (LCC) approach: A conceptual discussion of its usefulness for environmental decision-making. *Building and Environment*, 39(5):571–580, 5 2004. doi:10.1016/j.buildenv.2003.10.008.

[112] Danhong Wang, Jonas Landolt, Georgios Mavromatidis, Kristina Orehounig, and Jan Carmeliet. CESAR: A bottom-up building stock modelling tool for Switzerland to address sustainable energy transformation strategies. *Energy and Buildings*, 169:9–26, 6 2018. doi:10.1016/J.ENBUILD.2018.03.020.

[113] The UK government. A guide to energy performance certificates for the marketing, sale and let of dwellings Improving the energy efficiency of our buildings. Technical report, Department for Communities and Local Government, 2017. URL: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/671018/A_guide_to_energy_performance_certificates_for_the_marketing__sale_and_let_of_dwellings.pdf.

[114] Fazel Khayatian, Luca Sarto, and Giuliano Dall'O. Building energy retrofit index for policy making and decision support at regional and national scales. *Applied Energy*, 206:1062–1075, 11 2017. doi:10.1016/J.APENERGY.2017.08.237.

[115] Martin Achtnicht and Reinhard Madlener. Factors influencing German house owners' preferences on energy retrofits. *Energy Policy*, 68:254–263, 5 2014. doi:10.1016/J.ENPOL.2014.01.006.

[116] Constantine E. Kontokosta. Modeling the energy retrofit decision in commercial office buildings. *Energy and Buildings*, 131:1–20, 11 2016. doi:10.1016/J.ENBUILD.2016.08.062.

[117] Carl Christian Michelsen and Reinhard Madlener. Homeowners' preferences for adopting innovative residential heating systems: A discrete choice analysis for Germany. *Energy Economics*, 34(5):1271–1283, 9 2012. doi:10.1016/j.eneco.2012.06.009.

[118] Hassan Adan and Franz Fuerst. Modelling energy retrofit investments in the UK housing market: A microeconomic approach. *Smart and Sustainable Built Environment*, 4(3):251–267, 11 2015. doi:10.1108/SASBE-03-2013-0016.

[119] Baabak Ashuri, Hamed Kashani, and Jian Lu. An Investment Analysis Framework for Energy Retrofit in Existing Buildings. In *47th ASC Annual International Conference Proceedings*. Associated Schools of Construction, 2011.

[120] Graciela del Carmen Nava Guerrero, Gijsbert Korevaar, Helle Hvid Hansen, and Zofia Lukszo. Agent-Based Modeling of a Thermal Energy Transition in the Built Environment. *Energies*, 12(5):856, 3 2019. doi:10.3390/EN12050856.

[121] Xin Liang, Yi Peng, and Geoffrey Qiping Shen. A game theory based analysis of decision making for green retrofit under different occupancy types. *Journal of Cleaner Production*, 137:1300–1312, 11 2016. doi:10.1016/J.JCLEPRO.2016.07.200.

[122] William Swan, Les Ruddock, Luke Smith, and Richard Fitton. Adoption of sustainable retrofit in UK social housing. *Structural Survey*, 31(3):181–193, 7 2013. doi:10.1108/SS-12-2012-0039.

[123] Claudio Nägeli, Martin Jakob, Giacomo Catenazzi, and York Ostermeyer. Towards agent-based building stock modeling: Bottom-up modeling of long-term stock dynamics affecting the energy and climate impact of building stocks. *Energy and Buildings*, 211, 3 2020. doi:10.1016/J.ENBUILD.2020.109763.

[124] Xining Yang, Mingming Hu, Arnold Tukker, Chunbo Zhang, Tengfei Huo, and Bernhard Steubing. A bottom-up dynamic building stock model for residential energy transition: A case study for the Netherlands. *Applied Energy*, 306, 1 2022. doi:10.1016/J.APENERGY.2021.118060.

[125] Vincenzo Bianco and Clara Marmori. Modelling the deployment of energy efficiency measures for the residential sector. The case of Italy. *Sustainable Energy Technologies and Assessments*, 49:101777, 2 2022. doi:10.1016/J.SETA.2021.101777.

[126] Igor Sartori, Nina Holck Sandberg, and Helge Brattebø. Dynamic building stock modelling: General algorithm and exemplification for Norway. *Energy and Buildings*, 132:13–25, 11 2016. doi:10.1016/J.ENBUILD.2016.05.098.

[127] Horst Siebert. Environmental Quality as a Public Good. In *Economics of the Environment*, pages 59–95. Springer, Berlin, 2007. URL: https://link.springer.com/chapter/10.1007/978-3-540-73707-0_5.

[128] Kjell Arne Brekke, Snorre Kverndokk, and Karine Nyborg. An economic model of moral motivation. *Journal of Public Economics*, 87(9-10):1967–1983, 9 2003. doi:10.1016/S0047-2727(01)00222-5.

[129] R. Mark Isaac and James M. Walker. Group Size Effects in Public Goods Provision: The Voluntary Contributions Mechanism. *The Quarterly Journal of Economics*, 103(1):179, 2 1988. doi:10.2307/1882648.

[130] A Chaudhuri. Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature. *Exp. Econ.*, 14:47–83, 2011.

[131] Boyu Zhang, Xinmiao An, and Yali Dong. Conditional cooperator enhances institutional punishment in public goods game. *Applied Mathematics and Computation*, 390(125600), 2 2021. `doi:10.1016/j.amc.2020.125600`.

[132] Hal Varian. *Microeconomic Analysis.* Norton, New York, 3rd edition, 1992.

[133] Debra Israel and Arik Levinson. Willingness to Pay for Environmental Quality: Testable Empirical Implications of the Growth and Environment Literature. *Contributions in Economic Analysis & Policy*, 3(1):1–29, 2 2004. `doi:10.2202/1538-0645.1254`.

[134] Patricia A Champ, Kevin J Boyle, Thomas C Brown, and L George Peterson. *A primer on nonmarket valuation*, volume 3. Springer, 2003.

[135] Richard C. Bishop and Kevin J. Boyle. Reliability and Validity in Nonmarket Valuation. *Environmental and Resource Economics 2018 72:2*, 72(2):559–582, 1 2018. `doi:10.1007/S10640-017-0215-7`.

[136] Ian J. Bateman and Catherine L. Kling. Revealed Preference Methods for Nonmarket Valuation: An Introduction to Best Practices. *Review of Environmental Economics and Policy*, 14(2):240–259, 6 2020. `doi:10.1093/REEP/REAA009`.

[137] Klaus M. Miller, Reto Hofstetter, Harley Krohmer, and Z. John Zhang. How Should Consumers' Willingness to Pay be Measured? An Empirical Comparison of State-of-the-Art Approaches:. *Journal of Marketing Research*, 48(1):172–184, 2 2011. `doi:10.1509/JMKR.48.1.172`.

[138] James C. Anderson, Dipak C. Jain, and Pradeep K. Chintagunta. Customer Value Assessment in Business Markets: A State-of-Practice Study. *Journal of Business-to-Business Marketing*, 1(1):3–29, 6 1992. URL: `https://www.scholars.northwestern.edu/en/publications/customer-value-assessment-in-business-markets-a-state-of-practice-2`.

[139] R Inglehart, M Lagos, P Norris, E Ponarin, B Puranen, C Haerpfer, A Moreno, C Welzel, K Kizilova, J Diez-Medrano, M Lagos, E Ponarin, and B Puranen. World Values Survey: Round Six - Country- Pooled Datafile. Technical report, World Values Survey, 2014.

[140] Fredrik Carlsson and Peter Martinsson. Do Hypothetical and Actual Marginal Willingness to Pay Differ in Choice Experiments?: Application to the Valuation of the Environment. *Journal of Environmental Economics and Management*, 41(2):179–192, 3 2001. `doi:10.1006/JEEM.2000.1138`.

[141] Muzafer Sherif, Daniel Taub, and Carl I. Hovland. Assimilation and contrast effects of anchoring stimuli on judgments. *Journal of Experimental Psychology*, 55(2):150–155, 2 1958. doi:10.1037/H0048784.

[142] Adrian Furnham and Hua Chu Boo. A literature review of the anchoring effect. *The Journal of Socio-Economics*, 40(1):35–42, 2 2011. doi:10.1016/J.SOCEC.2010.10.008.

[143] Trudy A. Cameron and Michelle D. James. Estimating Willingness to Pay from Survey Data: An Alternative Pre-Test-Market Evaluation Procedure:. *Journal of Marketing Research*, 24(4):389–395, 12 1987. doi:10.1177/002224378702400406.

[144] A. Brett Hauber, Juan Marcos González, Catharina G.M. Groothuis-Oudshoorn, Thomas Prior, Deborah A. Marshall, Charles Cunningham, Maarten J. IJzerman, and John F.P. Bridges. Statistical Methods for the Analysis of Discrete Choice Experiments: A Report of the ISPOR Conjoint Analysis Good Research Practices Task Force. *Value in Health*, 19(4):300–315, 6 2016. doi:10.1016/j.jval.2016.04.004.

[145] George J. MacKerron, Catrin Egerton, Christopher Gaskell, Aimie Parpia, and Susana Mourato. Willingness to pay for carbon offset certification and co-benefits among (high-)flying young adults in the UK. *Energy Policy*, 37(4):1372–1381, 4 2009. doi:10.1016/j.enpol.2008.11.023.

[146] Lucia Rotaris, Marco Giansoldati, and Mariangela Scorrano. Are air travellers willing to pay for reducing or offsetting carbon emissions? Evidence from Italy. *Transportation Research Part A: Policy and Practice*, 142:71–84, 12 2020. doi:10.1016/j.tra.2020.10.014.

[147] Dalia Streimikiene, Tomas Balezentis, Ilona Alisauskaite-Seskiene, Gintare Stankuniene, and Zaneta Simanaviciene. A review of willingness to pay studies for climate change mitigation in the energy sector. *Energies*, 12(8), 4 2019. doi:10.3390/en12081481.

[148] Silvia Banfi, Mehdi Farsi, Massimo Filippini, and Martin Jakob. Willingness to pay for energy-saving measures in residential buildings. *Energy Economics*, 30(2):503–516, 2006. doi:10.1016/j.eneco.2006.06.001.

[149] Ecosystem Marketplace. Unlocking Carbon Market Potential In a Post-Paris World. Technical report, Ecosystem Marketplace, 2017. URL: https://www.ecosystemmarketplace.com/articles/unlocking-carbon-market-potential-in-a-post-paris-world/.

[150] Clinton J. Andrews, Daniel Yi, Uta Krogmann, Jennifer A. Senick, and Richard E. Wener. Designing buildings for real occupants: An agent-based approach. *IEEE Transactions on Systems, Man, and Cybernetics Part A:Systems and Humans*, 41(6):1077–1091, 2011. `doi:10.1109/TSMCA.2011.2116116`.

[151] Chris Silvia and Rachel M. Krause. Assessing the impact of policy interventions on the adoption of plug-in electric vehicles: An agent-based model. *Energy Policy*, 96:105–118, 2016. `doi:10.1016/j.enpol.2016.05.039`.

[152] J. R. Snape, P. J. Boait, and R. M. Rylatt. Will domestic consumers take up the renewable heat incentive? An analysis of the barriers to heat pump adoption using agent-based modelling. *Energy Policy*, 85:32–38, 2015. `doi:10.1016/j.enpol.2015.05.008`.

[153] J Richard Snape. *Incorporating human behaviour in an agent based model of technology adoption in the transition to a smart grid*. PhD thesis, DeMontfortUniversity, 2015.

[154] Varun Rai and Adam Douglas Henry. Agent-based modelling of consumer energy choices. *Nature Climate Change*, 6(6):556–562, 2016. `doi:10.1038/nclimate2967`.

[155] James Samuel Coleman, James Smoot Coleman, Thomas J. Farraro, and SAGE Publishing. *Rational Choice Theory: Advocacy and Critique*. SAGE Publications, Inc., California, USA, 1992. URL: `https://philpapers.org/rec/COLRCT`.

[156] Leo Breiman. Random Forests. *Machine Learning*, 45(1):5–32, 10 2001. `doi:10.1023/A:1010933404324`.

[157] Davide Chicco, Matthijs J. Warrens, and Giuseppe Jurman. The coefficient of determination R-squared is more informative than SMAPE, MAE, MAPE, MSE and RMSE in regression analysis evaluation. *PeerJ Computer Science*, 7:e623, 7 2021. `doi:10.7717/PEERJ-CS.623`.

[158] Mordecai Ezekiel. The Application of the Theory of Error to Multiple and Curvilinear Correlation. *Journal of the American Statistical Association*, 24(165):99, 3 1929. `doi:10.2307/2277015`.

[159] GRASS GIS. v.generalize - GRASS GIS manual, 2022. URL: `https://grass.osgeo.org/grass80/manuals/v.generalize.html`.

[160] David H Douglas and Thomas K Peucker. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *Cartographica: The International Journal for Geographic Information and Geovisualization*, 10(2):112–122, 10 2006. `doi:10.3138/FM57-6770-U75U-7727`.

[161] Met Office. Station history file for Watnall (including Hucknall) (DCnn: 4206, old number: 316/614, WMO number: 03354). Technical report, National Meteorological Library & Archive. URL: `https://library.metoffice.gov.uk/portal/Default/en-GB/RecordView/Index/248997`.

[162] M Eames and D Coley. On the creation of future probabilistic design weather years from UKCP09. *Building Services Engineering Research and Technology*, 32(2):127–142, 2011. `doi:10.1177/0143624410379934`.

[163] eppy 0.5.44 documentation. Welcome to eppy's documentation!, 2013. URL: `https://pythonhosted.org/eppy/`.

[164] Guangying Ren, Yeonsook Heo, and Minna Sunikka-Blank. Investigating an adequate level of modelling for retrofit decision-making: A case study of a British semi-detached house. *Journal of Building Engineering*, 26:100837, 11 2019. `doi:10.1016/j.jobe.2019.100837`.

[165] Tom De Smedt and Walter Daelemans. Pattern for Python Walter Daelemans. *Journal of Machine Learning Research*, 13:2063–2067, 2012. URL: `https://www.jmlr.org/papers/volume13/desmedt12a/desmedt12a.pdf`.

[166] Arun Kumar and Supriya P. Panda. A Survey: How Python Pitches in IT-World. *Proceedings of the International Conference on Machine Learning, Big Data, Cloud and Parallel Computing: Trends, Prespectives and Prospects, COMITCon 2019*, pages 248–251, 2 2019. `doi:10.1109/COMITCON.2019.8862251`.

[167] Shiva Verma. How Fast Numpy Really is and Why?, 2019. URL: `https://towardsdatascience.com/how-fast-numpy-really-is-e9111df44347`.

[168] The pandas development team. pandas-dev/pandas: Pandas 1.4.1, 2022. `doi:10.5281/zenodo.3509134`.

[169] S Gillies. The Shapely User Manual, 2022. URL: `https://shapely.readthedocs.io/en/stable/manual.html`.

[170] GeoPandas. GeoPandas 0.10.2+0.g04d377f.dirty, 2021. URL: `https://geopandas.org/`.

[171] QGIS. Welcome to the QGIS project!, 2022. URL: `https://qgis.org/en/site/`.

[172] Xiangwei Liu, Youming Chen, Hua Ge, Paul Fazio, and Guojie Chen. Determination of Optimum Insulation Thickness of Exterior Wall with Moisture Transfer in Hot Summer and Cold Winter Zone of China. *Procedia Engineering*, 121:1008–1015, 1 2015. `doi:10.1016/J.PROENG.2015.09.072`.

[173] António M. Raimundo, Nuno B. Saraiva, and A. Virgílio M. Oliveira. Thermal insulation cost optimality of opaque constructive solutions of buildings under Portuguese temperate climate. *Building and Environment*, 182:107107, 9 2020. `doi:10.1016/J.BUILDENV.2020.107107`.

[174] Robert Dylewski. Optimal Thermal Insulation Thicknesses of External Walls Based on Economic and Ecological Heating Cost. *Energies*, 12(18):3415, 9 2019. `doi:10.3390/EN12183415`.

[175] Claire Smith. English Housing Survey Introduction and main findings. Technical report, Ministry of Housing Communities and Local Government, 2020. URL: `https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1000108/EHS_19-20_Energy_report.pdf`.

[176] Robert Parsons, editor. *Fundamentals: 2005 Ashrae Handbook*, volume 1. American Society of Heating, Refrigerating and Air-Conditioning Engineers, 2005.

[177] Geoff Hammond and Craig Jones. INVENTORY OF CARBON & ENERGY (ICE) Version 1.6a. Technical report, University of Bath, Bath, 2008. URL: `https://perigordvacance.typepad.com/files/inventoryofcarbonandenergy.pdf`.

[178] I Boustead. Eco-profiles of the European Plastics Industry. Technical report, PlasticsEurope, 2005. URL: `www.plasticseurope.org.`

[179] M. D. McKay, R. J. Beckman, and W. J. Conover. A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code. *Technometrics*, 21(2):239, 5 1979. `doi:10.2307/1268522`.

[180] Odysee Project. Sectoral Profile - Households Energy consumption. Technical report, Odyssee-Mure, 2019.

[181] The UK government. Report 4: Main heating systems. Technical report, Department of Energy and Climate Change, 2013. URL: `https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/274772/4_Main_heating_systems.pdf`.

[182] Diederik P. Kingma and Jimmy Lei Ba. Adam: A Method for Stochastic Optimization. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, 12 2014. URL: `https://arxiv.org/abs/1412.6980v9`.

[183] Cameron R. Lawrence, Russell Richman, Maria Kordjamshidi, and Cecilia Skarupa. Application of surrogate modelling to improve the thermal performance of single-family homes through archetype development. *Energy and Buildings*, 237:110812, 4 2021. `doi:10.1016/J.ENBUILD.2021.110812`.

[184] Ehsan Asadi, Manuel Gameiro Da Silva, Carlos Henggeler Antunes, Luís Dias, and Leon Glicksman. Multi-objective optimization for building retrofit: A model using genetic algorithm and artificial neural network and an application. *Energy and Buildings*, 81:444–456, 2014. doi:10.1016/j.enbuild.2014.06.009.

[185] EWI Store. Home, 2022. URL: https://ewistore.co.uk/.

[186] Cut Price Insulation. Quinn Therm PIR Insulation 1200 x 2400 x 25mm, 2022. URL: https://www.cutpriceinsulation.co.uk/products/quinn-therm-pir-insulation-1200-x-2400-x-25mm?variant=16705205829689&currency=GBP&utm_source=google&utm_medium=cpc&utm_campaign=google+shopping&gclid=CjwKCAjwp-X0BRAFEiwAheRui3xgQlBY8_g-rYFJwCGghWACoqK9Lx1ExhnHGiTgEekWY7oTEHIQAxoC2xwQAvD_BwE.

[187] Homebuilding. Retrofit Central Heating, 2020. URL: https://www.homebuilding.co.uk/advice/retrofit-central-heating.

[188] GetAWindow. Double Glazing Cost: All The Information You Need, 2022. URL: https://www.getawindow.co.uk/window-glass/double-glazing-cost/.

[189] Jeremy J. Siegel. The real rate of interest from 1800-1990. A study of the U.S. and the U.K. *Journal of Monetary Economics*, 29(2):227–252, 4 1992. doi:10.1016/0304-3932(92)90014-S.

[190] Lukasz Rachel and Thomas D Smith. Are Low Real Interest Rates Here to Stay? *. *International Journal of Central Banking*, 50, 8 2017.

[191] BREGroup. Standard Assessment Procedure (SAP 2012), 2012. URL: https://www.bregroup.com/sap/standard-assessment-procedure-sap-2012/?cn-reloaded=1.

[192] Jonathan Reynolds, Yacine Rezgui, Alan Kwan, and Solène Piriou. A zone-level, building energy optimisation combining an artificial neural network, a genetic algorithm, and model predictive control. *Energy*, 151:729–739, 2018. doi:10.1016/j.energy.2018.03.113.

[193] Ali Alajmi and Jonathan Wright. Selecting the most efficient genetic algorithm sets in solving unconstrained building optimization problem. *International Journal of Sustainable Built Environment*, 3(1):18–26, 6 2014. doi:10.1016/J.IJSBE.2014.07.003.

[194] Tiejun Li, Guifang Shao, Wangda Zuo, and Sen Huang. Genetic algorithm for building optimization - State-of-the-art survey. *ACM International Conference Proceeding Series*, Part F128357:205–210, 2 2017. doi:10.1145/3055635.3056591.

[195] Nancy Chinchor. MUC-4 evaluation metrics. *4th Message Understanding Conference, MUC 1992 - Proceedings*, pages 22–29, 1992. doi:10.3115/1072064.1072067.

[196] Tom O'Malley, Elie Bursztein, James Long, François Chollet, Haifeng Jin, Luca Invernizzi, and others. KerasTuner. https://github.com/keras-team/keras-tuner, 2019. URL: https://github.com/keras-team/keras-tuner.

[197] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization. *Journal of Machine Learning Research*, 18:1–52, 3 2016. URL: https://arxiv.org/abs/1603.06560v4.

[198] Neil Higgins. Dwelling Stock Estimates: 2017, England. Technical report, Office for National Statistics, 2018. URL: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/710382/Dwelling_Stock_Estimates_2017_England.pdf.

[199] T. Ibn-Mohammed, R. Greenough, S. Taylor, L. Ozawa-Meida, and A. Acquaye. Operational vs. embodied emissions in buildings—A review of current trends. *Energy and Buildings*, 66:232–245, 11 2013. doi:10.1016/J.ENBUILD.2013.07.026.

[200] The UK government. Calculate the carbon dioxide equivalent quantity of an F gas, 2014. URL: https://www.gov.uk/guidance/calculate-the-carbon-dioxide-equivalent-quantity-of-an-f-gas.

[201] US Environmental Protection Agency. Greenhouse Gas Equivalencies Calculator, 2022. URL: https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator.

[202] European Environment Agency. carbon dioxide equivalent, 2022. URL: https://www.eea.europa.eu/help/glossary/eea-glossary/carbon-dioxide-equivalent.

[203] Stig-Inge -I Gustafsson and Björn G. Karlsson. Why is life-cycle costing important when retrofitting buildings. *International Journal of Energy Research*, 12(2):233–242, 1988. doi:10.1002/er.4440120206.

[204] Xing Su, Shaochen Tian, Xiaolu Shao, and Xuan Zhao. Embodied and operational energy and carbon emissions of passive building in HSCW zone in China: A case study. *Energy and Buildings*, 222:110090, 9 2020. doi:10.1016/j.enbuild.2020.110090.

[205] Parliament of the United Kingdom. Climate Change Act 2008. Technical report, 2008.

[206] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical Bayesian Optimization of Machine Learning Algorithms. In F Pereira, C J C Burges, L Bottou, and K Q Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. URL: `https://proceedings.neurips.cc/paper/2012/file/05311655a15b75fab86956663e1819cd-Paper.pdf`.

[207] Scott M. Lundberg and Su In Lee. A Unified Approach to Interpreting Model Predictions. *Advances in Neural Information Processing Systems*, 2017-December:4766–4775, 5 2017. URL: `https://arxiv.org/abs/1705.07874v2`.

[208] S Mazzanti. SHAP Values Explained Exactly How You Wished Someone Explained to You, 2020. URL: `https://towardsdatascience.com/shap-explained-the-way-i-wish-someone-explained-it-to-me-ab81cc69ef30`.

[209] S Lundberg. shap/overviews.rst at master · slundberg/shap, 2020. URL: `https://github.com/slundberg/shap/blob/master/docs/overviews.rst`.

[210] Claudia Keser and Frans Van Winden. Conditional Cooperation and Voluntary Contributions to Public Goods. *Scandinavian Journal of Economics*, 102(1):23–39, 3 2000. `doi:10.1111/1467-9442.00182`.

[211] Dustin Tingley and Michael Tomz. Conditional Cooperation and Climate Change. *Comparative Political Studies*, 47(3):344–368, 11 2013. `doi:10.1177/0010414013509571`.

[212] Eli Cohen. Applying best-worst scaling to wine marketing. *International Journal of Wine Business Research*, 21(1):8–23, 3 2009. `doi:10.1108/17511060910948008`.

[213] T N Flynn and A A J Marley. Best Worst Scaling: Theory and Methods Handbook of Choice Modelling. In *Handbook of Choice Modelling*, pages 1–29. 2007.

[214] Jordan Louviere, Ian Lings, Towhidul Islam, Siegfried Gudergan, and Terry Flynn. An introduction to the application of (case 1) best-worst scaling in marketing research. 2013. `doi:10.1016/j.ijresmar.2012.10.002`.

[215] Pat Auger, Timothy M. Devinney, and Jordan J. Louviere. Using best-worst scaling methodology to investigate consumer ethical beliefs across countries. *Journal of Business Ethics*, 70(3):299–326, 2 2007. `doi:10.1007/S10551-006-9112-7`.

[216] Nottingham City Council. Planning Applications in Nottingham, 2021. URL: `https://www.opendatanottingham.org.uk/dataset.aspx?id=28`.

[217] D White. How often do we move house in Britain? - Zoopla, 2017. URL: https://www.zoopla.co.uk/discover/property-news/how-often-do-we-move-house-in-britain/.

[218] Jordan J. Louviere, Terry N. Flynn, and A. A.J. Marley. *Best-Worst Scaling: Theory, Methods and Applications*. Cambridge University Press, 2015. doi:10.1017/CBO9781107337855.

[219] Nicolas Krucien, Verity Watson, and Mandy Ryan. Is Best–Worst Scaling Suitable for Health State Valuation? A Comparison with Discrete Choice Experiments. *Health Economics*, 26(12):e1–e16, 12 2017. doi:10.1002/HEC.3459.

[220] Liz Waters. Energy Consumption in the UK (ECUK) 1970 to 2019. Technical Report July 2019, Office of National Statistics, 2019. URL: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1020152/2020_Energy_Consumption_in_the_UK__ECUK_.pdf.