INTERNATIONAL BUILDING PERFORMANCE SIMULATION ASSOCIATION ENGLAND

Proceedings of BSO-V 2020
5th Building Simulation and Optimization Virtual Conference
21-22 September 2020
Loughborough, UK

Surrogate Optimisation of Housing Stock Retrofits using Deep Neural Networks

James Hey[1], Peer-Olaf Siebers[1], Paul Nathanail[2], Ender Özcan[1], Darren Robinson[3]
[1]School of Computer Science, University Of Nottingham, United Kingdom
[2] Land Quality Management Ltd, Nottingham, UK
[3]School of Architecture, University of Sheffield, United Kingdom

Abstract

Surrogate modelling can greatly reduce the computational time required to perform building simulation, by trading accuracy for execution speed. We propose a surrogate optimisation extension to this process, capturing the wider optimisation loop in a second surrogate model to predict the optimisation output from runs of a conventional surrogate model; applying this to optimise of building energy retrofit strategies. In this we model the housing stock of part of Nottingham (UK), representing the c. 95,000 dwellings using a combination of Ordnance Survey and English Housing Survey data.

We use established simulation and optimisation methods to create a sample of 5000 near-optimal retrofit solutions for buildings in Nottingham. Using this sample we train a set of DNN models to form a Surrogate Optimiser to predict retrofits for the remaining building stock. Using this method, a cost efficient whole house retrofit solution was found in 16.7% of the housing stock, compared with 19.2% identified by the base optimiser. The solutions identified by the optimiser scored 11% worse than those identified by the base optimiser, but the surrogate optimiser was approximately 100,000 times faster; although this improvement drops to c.20x when considering the time required to generate the training data.

Introduction

In 2018, domestic dwellings accounted for 35% (414 TWh) of the UK's total energy useage, mainly (65%) for space heating using natural gas boilers or electric heaters (Office for National Statistics, 2019). Achieving the UK government's target of net zero greenhouse gas emissions by 2050 [1], will require widespread retrofit of the existing housing stock.

Simulation is a tool which allows researchers and industry professionals to evaluate the impact of changes to a system before implementing them in the real world, allowing assessments of the costs and benefits before incurring them. The integration of simulation with optimisation provides for a powerful decision support tool. However, these optimisation methods can entail performance challenges, as a large number of iterations may be required to converge on a near optimal solution. For some simulations this process is prohibitively resource expensive. To overcome this challenge, surrogate modelling has been used: machine learning models which capture the behaviour of the simulator, replacing the simulation at optimisation run time, yielding significant performance increases.

We consider a situation in which optimisation is desired for a large real-world housing stock, such as producing an optimal building energy retrofit for the entire domestic stock of (part of) a city. This could be used to identify the most beneficial retrofits for stakeholders such as social housing associations, councils or energy service companies. But even using a surrogate modelling approach, this may be resource prohibitive. To address this issue, we have developed a framework for creating a Surrogate Optimiser, in which the optimisation phase is incorporated into a higher level surrogate. This extends established methods by creating a pre-trained general model using a sample of the building stock. This model allows rapid optimisation of the remaining stock.

To this end, we employ dynamic building energy simulation software to generate a training set for a surrogate energy demand model, which is used by a genetic algorithm to optimise a sample of a city's housing stock. This sample is then used to train a set of Deep Neural Networks (DNNs) to form a Surrogate Optimiser that is able to output near-optimal retrofit solutions for buildings of varying geometry, location and initial building fabric states. This is a first attempt at developing a Surrogate Optimiser of retrofit strategies, and as such we made simplifications regarding building occupancy, heating schedules and stakeholders' objective functions, which were fixed to contain the dimensionality of the input space.

Background and Related Works

Building Energy Simulation

There exist several packages for building energy simulation, but the open source software EnergyPlus (Crawley et al., 2000) has achieved a dominant position. EnergyPlus is primarily written in C++, al-

[1]http://www.legislation.gov.uk/uksi/2019/1056

lowing for relatively high speed simulations; although the complexity of these simulations means that run times are still prohibitive to perform at scale. EnergyPlus takes an input file with defined parameters and after running the simulation generates a set of output files. This direct input to output mapping grants potential for function space mapping of the subsystem using machine learning techniques, given a sufficiently large sample size.

Surrogate Models for Optimisation

Sometimes referred to as emulators or response surface models, Surrogate Models are meta-models which attempt to substitute a complex and computationally expensive simulation process with a faster model that emulates the simulation it is replacing(Eisenhower et al., 2012). This is particularly useful when performing optimisation, as this tends to require a very large number of iterations which may be infeasible if the simulation model is run at each iteration. Surrogate models attempt to create a function that maps a simulation system's input space to the output space without requiring intensive run time simulation (Asadi et al., 2014).

The importance of using a surrogate model for this process lies in speed. The optimisation process is slow, for example (Aijazi, 2017) note 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.

Prada et al. (2018) provide an analysis of the performance of different surrogate modelling techniques. They used a genetic algorithm with surrogate models and compared results to a brute force optimisation to determine the efficacy of different 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. Of the modelling techniques tested, they found the Multivariate adaptive regression spline (MARS) to be more efficient.

Wate et al. (2020) demonstrate a novel surrogate model framework which emulates a stochastic building performance simulator, trained using EnergyPlus co-simulated with the Multi-Agent Stochastic Simulation platform NoMASS to simulate occupants' behaviours. The authors demonstrate 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.

The above surrogate models for building energy modelling are used for investigative purposes rather than for optimisation. In the present context, optimisation of building energy performance involves the selection of building properties which reduce energy usage in accordance with some function (cost, environmental impact etc). Retrofitting is the process of adjusting elements of an existing building to make it more energy efficient. Aijazi (2017) used surrogate modelling to perform retrofit optimisation. This was done by first extracting 11 input features from the building, such as wall and roof insulation, heating system efficiency and window to wall ratio. Where input variables are correlated, uncorrelated linear transformations of input variables are used (Aijazi, 2017, p. 41)[2]. In their previous work, Asadi et al. (2014) used an Artificial Neural Network (ANN) as their surrogate model[3]. The ANN was trained using a building energy performance simulator, with input neurons representing the categories of retrofit intervention possible, and the output representing physical outputs such as energy required for heating and cooling the building. The input space was comprised of 5 categories, ranging in size from 4 to 24 possible parameter values, giving 20,736 combinations of retrofit. The ANN was trained using 950 of these, approximately 5% of the input space. The neural network proved to be effective, with average errors of 2.5% or less for all output variables when compared with simulated test data. A Genetic Algorithm (GA) was used to create candidate solutions during optimisation, which were evaluated with the neural network, the outputs of which were fed into a multi-criteria decision module which evaluated the trade-off between cost, energy saving and comfort improvements.

Asadi et al. (2014) also utilise a surrogate model, referred to as a Response Surface Approximation model, to perform retrofit optimisation on a school building. After creating a simulation model, they used Latin Hypercube Sampling (LHS) to generate a sample of 950 possible retrofits to train an artificial neural network (ANN). They placed this trained ANN inside a multi-objective genetic algorithm to perform optimisation. Retrofits included in the model ranged from wall and window insulation, to HVAC system choice and even a solar collector, with model outputs of cost, thermal comfort and energy consumption.

Finally, Reynolds et al. (2018) create a predictive ANN to use as a surrogate model for evaluation in a Genetic Algorithm (GA). Their work focused on

[2]It is worth noting here that these features are relatively crude metrics for capturing geometric aspects of the data. In a city which may have a wide range of geometric configurations a model trained with limited geometric parameters could perform poorly even with a representative sample set.

[3]The authors did not use the term surrogate model or any of it's common synonyms(emulator, space map, response surface model etc.), but the ANN used did emulate the behaviour of a building energy model, which it was trained with.

the thermostat control of 6 control zones in an office building in Cardiff. Their optimisations focused on reducing the energy consumption (or cost when modelled as a two tariff system).

The above papers successfully demonstrate the application of surrogate modelling techniques to support energy simulation and, in conjunction with conventional techniques, the optimisation of building energy retrofit decisions. However these have been for single buildings having a specific geometry, and typically for relatively simple sets of interventions, hand-tailored by the experimenters. With a significant amount of both computational and human labour required per building, the solutions are not scalable to large stocks of building, or general in nature.

Model Based Methods for Optimisation

Model based optimisation methods exist 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 (Costa et al., 2018). These methods are quite promising and have been proposed as alternatives to more common metaheuristic approaches such as GA optimisation (Wortmann, 2018). While deploying a similar concept to the method proposed, these algorithms are trained in real time to optimise a specific objective function, in contrast to pre-trained model to achieve fast optimisations on genetic buildings from a data set by training using a metaheuristic approach.

Building Stock Models

Sousa et al. (2017) reviewed 29 housing stock energy models (HSEM) and found that the tools being utilised by researchers are "limited in their scope and employ simplistic models that limit their utility". One criticism leveled at these HSEMs was their rigidity, lacking modular design that can account for household investment decisions, a criticism which was central in the modular design of our proposed method, which is laid out in the following section.

Melo et al. (2014) performed a feasibility study examining the use of an ANN for surrogate modelling the urban scale building stock of Florianópolis, Brazil. They built and analysed a surrogate model using an ANN, by representing the building stock using 16 typologies and simulating 200 variants of each. Using these 3600 samples they performed sensitivity analysis and settled on a single hidden neuron layer of 9 neurons. The ANN performed well and their analysis demonstrated that a range of physical properties were needed to construct the model (with no single attribute dominating performance results). They also conclude that users of linear regression models for building energy simulation should "replace these models with ANN models, due to the simplicity and improved flexibility and accuracy of ANNs."(Melo et al., 2014, p.465).

Method

The process of developing a Surrogate Optimiser begins with development of a surrogate model using the basic input-output mappings generated by standard simulation. This process is then extended by encapsulating this surrogate modelling technique within a broader machine learning, as shown in Figure 1. This model is trained with mappings from the initial candidate solutions directly to the near-optimal solutions generated by the optimisation loop.

One difficulty that may arise with a Surrogate Optimiser is the inflexibility of mapping an initial case to a single near optimal solution. This isn't a problem in the case of a single objective, or where an objective function can be designed a-priori to synthesise multiple factors to compare potential solutions. But in cases where we wish to consider multiple objectives, such as the cost of an improvement and non-financial benefits, a single objective method will not suffice. In these instances we suggest the production of a Pareto front of solutions. A Pareto-optimal solution is one which cannot be dominated by another solution across all of the factors of interest, requiring other Pareto solutions to trade off in at least one factor to gain in the others. For example in an abstract two factor optimisation that considers both the utility and cost of a decision, a set of Pareto-optimal solutions will vary from those with low cost, low utility to high cost and high utility, but no solution will be both lower cost and higher utility than another solution in the Pareto set.

In the following Section, we discuss in detail the application of the framework summarised in Figure 1.

Case Study

To demonstrate the efficiency gains of Surrogate Optimisation, we study the domestic building stock of the city of Nottingham, UK. A data set of 95,500 domestic dwellings was used for sampling. With the exception of blocks of flats[4], this represents the entire housing stock. This data set was created using a combination of Ordinance Survey (OS) data, statistical distribution of attributes based on the English Housing Survey (EHS) and the 2011 census. As such, factors like building footprint, and eave and ridge height are accurate to the level of the individual building. Other attributes such as existing envelope features (levels of insulation, glazing etc.) were either extracted from Energy Performance Certificate (EPC) data or disaggregated from regional EHS data where building specific data was unavailable.

In our optimisations, we consider a stakeholder who wishes to identify which buildings in the stock have the greatest retrofit potential. This could, for example, be a retrofit installer looking to target buildings that are most likely to benefit from improvements

[4] these were excluded, as the data available was insufficiently granular to reliably model them

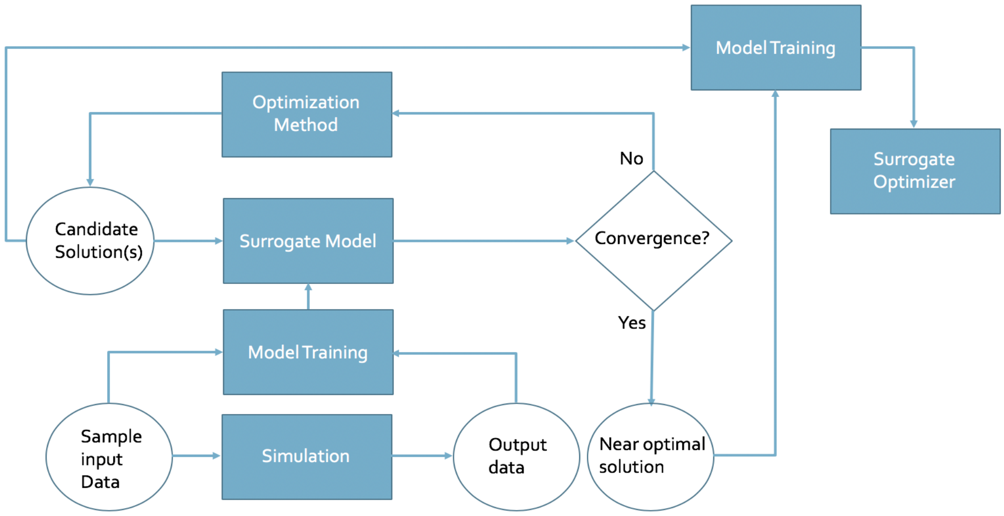

Figure 1: Optimisation framework using surrogate model

in the building fabric or from system substitution. All experiments are conducted on a 12 core Ryzen 9 3900x 3.80GHz Linux machine with 16.00GB memory.

Energy Performance Surrogate Model

The first stage of developing the Surrogate Optimiser requires the generation of near-optimal retrofits. To do this we first determine the desired set of possible retrofits to consider, simulate a sample of buildings with a variety of retrofits, then train a surrogate energy model using these simulation results.

Determining Retrofit Genes

In order to optimise using a Genetic Algorithm (GA), our possible retrofit solutions need to be expressible as a combination of genes. These were made deliberately simple for the purposes of this proof-of-principle study. A candidate solution encodes eight retrofit options. These include internal and external wall insulation of three different materials with eight possible thicknesses, varying between 30 and 100 millimeters. Windows can be retrofit with single (though this is admittedly unlikely), double or triple glazing. Only a single material (mineral wool) was considered for roof insulation, but possible thicknesses range from 50mm to 300mm, in 25mm increments. Finally, the option of retrofitting gas central heating is considered. This was included to create a potentially interesting trade-off between cost savings and carbon reduction (since the cost of gas is significantly less, but cannot be de-carbonised in the way the grid can be). The raw number of combinations of possible retrofit settings is 110,592 for each and every building. However, this will be lower in practice, as retrofit decisions will depend on the building's initial state (e.g. a building which already has triple glazing would disregard the glazing replacement option). A full description of the possible retrofit settings is provided in Table 1.

Simulation of retrofits

To train the surrogate energy model, a broad range of input and output data was required. Simulations were performed using EnergyPlus, and so an EnergyPlus input data file (IDF) template was constructed [5], with semantic attributes drawn from survey data to generate a set of input combinations as shown in Figure 2. As such, the simulation of aggregate energy use is expected to be plausible, but building-specific results may not be (in the absence of building-specific survey data).

Training the Energy Performance Surrogate Model

To train our Surrogate Model, a data set of 10,000 building states was generated. The base buildings that the hypothetical retrofits were applied to were selected using Latin Hypercube Sampling over the building's floor area, age and form. The selected base buildings were then mutated, with random genes assigned to each one. This created a data set with base buildings representative of the building stock, and a wide range of gene values across that sample. The heating demand of these hypothetical building states was simulated, with these simulations taking an average of 6.3 CPU seconds to complete.

The model was trained with a Deep Neural Network (DNN). First attempts at model creation used single layer ANNs, but these models struggled to perform well for the wide variety of buildings found in the stock. The structure of the DNN was determined through trial and error, guided by the Mean Absolute Percentage Error (MAPE) and r^2 values. The selected training structure contained 5 hidden layers, with a diminishing number of neurons per layer, from the input layer of 20 to the final hidden layer of 12 neurons. The total number of hidden neurons was 67.

[5] in a similar manner to that used by the dynamic housing stock energy simulation platform *EnHub* (Sousa et al., 2018)

Component	Gene Name	Possible Values	Unit
External Wall Insulation Material	EWI_mat	None (uninsulated), XPS, EPS, PIR	N/A
External Wall Insulation Thickness	EWI_thick	0.03-0.1 in increments of 0.01	Meters
Internal Wall Insulation Material	IWI_mat	None (uninsulated), XPS, EPS, PIR	N/A
External Wall Insulation Thickness	IWI_thick	0.03-0.1 in increments of 0.01	Meters
Heating Method	Heating	Electric, Gas Central Heating	N/A
Roof Insulation Material	Roof_Mat	None, Mineral Wool	Boolean
Roof Insulation Thickness	Roof_thick	0.05-0.3 in increments of 0.025	Meters
Glazing Type	Glazing	Single Glazing, Double Glazing, Triple Glazing	N/A

Table 1: Description of possible retrofit installation

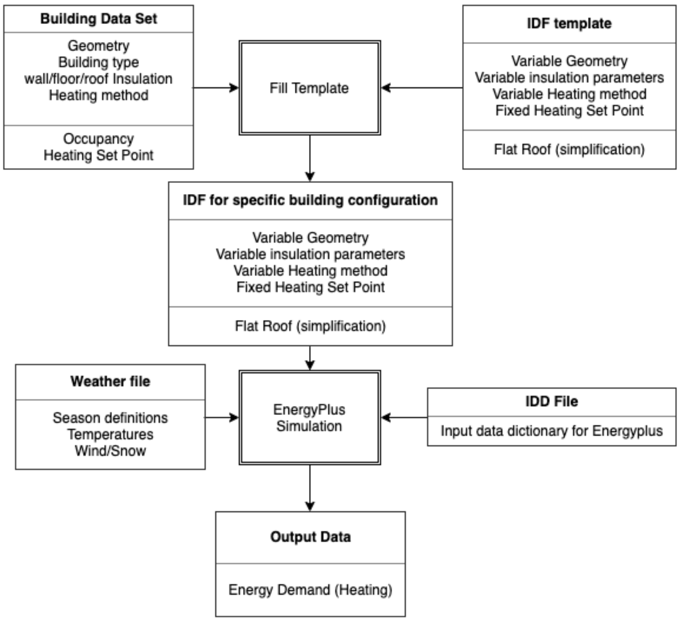

Figure 2: Process flow for training the surrogate energy model

The network was initialised with random weights, and was trained using a batch size of 5 for 100 epochs. The model was trained using 70% of the data, with the remaining 30% used for model testing. The model predicted the test data with a MAPE of 3.8% and an r^2 of 0.978.

Optimisation using Genetic Algorithm

A GA was used to obtain the required set of near-optimal solutions. The process flow for obtaining these solutions is demonstrated in Figure 3

Costing Model

The installation cost of retrofits was calculated based on the intervention chosen. The 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). Installation costs are also averages, which attempt to capture additional costs such as reparation of thermal bridges created from the intervention. This costing component is designed to be modular, allowing for more granular costing models if they are required for a certain use case. Full retrofit costing tables are laid out in Table 2.

Objective Function

To score candidate solutions, an objective function was required. The objective function utilised was the Net Present Value (NPV) of the savings made over a 25 year period

Genetic Algorithm Settings

Using this objective function, a simple evolutionary algorithm spawned a pool (population) of candidate solutions (individuals), which were scored using the energy demand calculated by the trained Surrogate Model as the objective/fitness function. The randomly created initial population of solutions is evolved making use of the genetic operators of crossover and mutation. At each evolutionary cycle, a new population of solutions is created which then replaces the old generation. A tournament selection

Material Gene Name	Gene Value	Material Name	Min Cost	Max Cost	Notes
EWI_mat	1	XPS	109.50	122.88	£ per squere meter
EWI_mat	2	EPS	106.67	115.00	£ per squere meter
EWI_mat	3	PIR	107.99	115.58	£ per squere meter
IWI_mat	1	XPS	96.67	105.00	£ per squere meter
IWI_mat	2	EPS	97.99	105.58	£ per squere meter
IWI_mat	3	PIR	97.99	105.58	£ per squere meter
Roof_mat	1	MineralWool	33.26	46.09	£ per squere meter
Heating	1	Central heating	3,000.00	3,000.00	£ per installation
Glazing	0	Single Glazing	-	-	Excluded (dominated)
Glazing	1	Double Glazing	375.00	375.00	£ per window
Glazing	2	Triple Glazing	475.00	475.00	£ per window

Table 2: Cost of retrofit installations

operator is used to choose two parent solutions that will undergo crossover. First, a number of individuals (tournament size) are selected randomly, then one of those solutions is selected as a parent based on a tournament pressure favouring the fitter individuals (solutions with higher quality). The same process is repeated for the second parent. Then two children/new solutions are created by applying the traditional one-point crossover operator to the chosen parents. The mutation operator is applied to both of the new solutions in a standard manner with the mutation rate of 5%. Mutations were directed away from dominated spaces to speed up convergence (for example, replacing existing insulation with less thick substitutes was not considered). A pool of half of the population size of new solutions were created by repeating this process and applying the genetic operators. Then those new solutions replaced the randomly selected individuals from the current population based on roulette wheel method. The evolutionary process was terminated when the maximum number of generations was reached, or when the average percentage change of the best score over the last 5 (stall) generations was sufficiently small (tolerance). GA returns the best solution in the final population. A full report of GA settings can be found in Table 3.

GA Parameter	Setting
Number of Genes	8
Population Size	500
Initialisation	Random
Selection Operator	Tournament
Tournament Size	20
Tournament Pressure	0.9
Crossover Operator	One-point
Mutation Operator	Traditional
Mutation Rate	0.05
Replacement Operator	Roulette Wheel
Tolerance	0.0005
Stall Generations	5
Max. Number of Generations	200

Table 3: Genetic Algorithm Settings

Genetic Algorithm Results

Due to the relatively small search space, the majority of optimisations converged before reaching the maximum generations, resulting in an average optimisation time of 32.3 CPU seconds per building. The GA was run on a sample of 5264 buildings from the Nottingham housing stock. Material, labour and electricity prices were set to realistic present day levels and a discount rate of 2% was used to calculate the present value of savings. With these parameters, retrofits were found to be profitable in 1,011 (19.2%) cases.

Surrogate Optimiser Training

The optimisation output generates a near-optimal retrofit per building optimised; this optimised state being highly dependent on the input building's initial state. To create a predictive Surrogate Optimiser (shown in Figure 4) we used a set of 8 DNNs, one to predict the optimal value of each gene. The sample of 1,011 profitable retrofits discussed above was used as an initial training set. In order to discourage the network from echoing input genes as optimal, the 4,253 buildings for which no positively scoring retrofit was found were removed from the sample set.

The performance of the neural networks naturally varied, as some optimisations were easier to predict than others. For example, internal wall insulation (IWI) dominated external wall insulation (EWI) in this scenario, because it costs less while providing similar thermal benefits. As such the IWI material predictor performed quite well as a classifier with a weighted f-1 value of 0.89, with the thickness predictor obtaining an f-1 of 0.81.

Surrogate Optimiser Results

Scoring the trained Optimiser using traditional classification metrics is less appropriate in this case, since the score of the combined gene output represents the desired output of the Optimiser. Luckily, the optimised gene outputs are simple to score using the same NPV of savings metric discussed when performing the original optimisation. The optimiser was used to generate optimisations for the whole city data set of

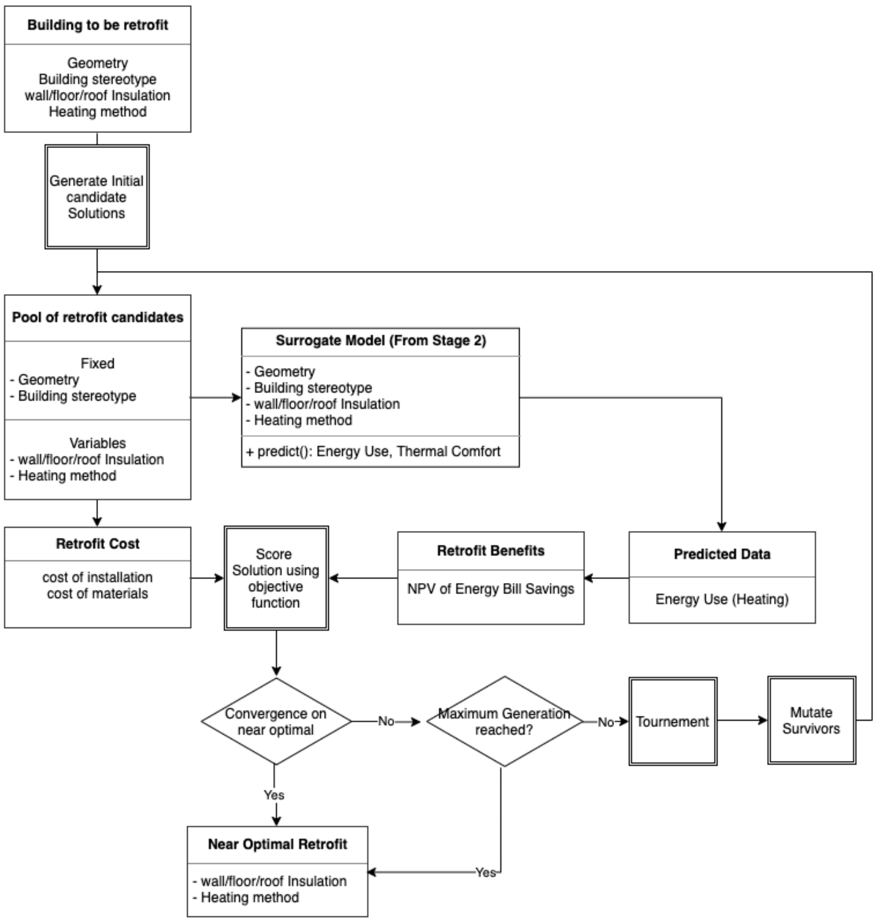

Figure 3: Process flow of retrofit optimisation process

95,500 buildings. This prediction took 36.5 CPU seconds, approximately the same time required to optimise a single building using the Surrogate Model embedded into the GA. The Surrogate Optimiser identified a profitable retrofit for 16.7% of the building stock. This compares with 19.2% found in the sample of buildings optimised using the GA. This suggests the Optimiser was able to identify 87% of the buildings for which a profitable optimisation could be found using a GA. The average score of positively scoring retrofits identified by the GA was £1,816.68, while the equivalent metric for the Surrogate Optimiser was £1,615.82, making the Surrogate Optimiser predictions 11% worse (by objective value) than the solutions generated by the GA. On the other hand, the pre-trained optimiser was able to predict the optimised scores approximately 100,000 times faster than the GA.

Discussion and Future Work

The results of the case study demonstrated the potential for surrogate optimisation to predict near optimal solutions on large data sets made up of non-trivial but relatively well behaved optimisation problems; problems that lend themselves well to a DNN mapping. It is likely that if more chaotic elements were introduced, such as non rational human behaviour, the stable island of optimisation vectors would break down, providing a greater challenge in training a surrogate optimiser. Even within the relatively constrained gene choices used to test this concept, there was an 11% loss in accuracy from moving from a traditional GA Optimiser to a Surrogate Optimiser. However, the speed savings were significant. With approximately 5% of the entire building stock used to train the optimiser, it was 20 times faster. This suggests that it is feasible to use a Surrogate Optimiser to evaluate retrofit options at the scale of an entire building stock, for example to support building stock decarbonisation policy measures (e.g raising carbon tax), the effects of future energy price changes or technology changes.

Given the trend towards massive parallelisation of computation, especially for problems like building stock modelling (which is highly parallelisable), there may be some doubt as to the utility of a pre-trained Surrogate Optimiser. In the case study chosen, for example, the computational cost of optimising the entire population is technically feasible. However, an optimiser of this nature may become a necessity when many building stock optimisation stages are required, such as optimisation of a proposed carbon tax policy.

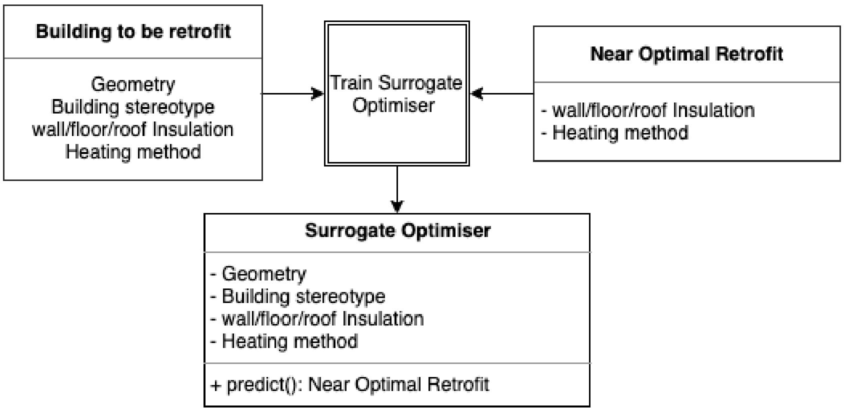

Figure 4: Process flow for training the Surrogate Optimiser

Additionally, the data generation stage of producing the Surrogate Optimiser is similarly parallelisable, so the speed improvements obtained can be retained for distributed systems.

There is considerable scope to further develop this proof-of-principle Surrogate Optimiser. This could include the expansion of the gene choice to encompass a broader or more challenging range of retrofits, the integration of exogenous scenarios (e.g. future energy costs) into the model training phase, or a more detailed analysis of the modelling techniques employed by the optimiser to close the outcome gap between Surrogate Optimisation and traditional retrofit optimisation methods. For example the DNN hyperparameters, GA settings, or input data processing (including the size of the training data set) could be further tuned to attempt to improve outcomes.

References

Aijazi, A. N. (2017). *Machine Learning Paradigms for Building Energy Performance Simulations.* Ph. D. thesis, Massachusetts Institute of Technology.

Asadi, E., M. G. D. Silva, C. H. Antunes, L. Dias, and L. Glicksman (2014). Multi-objective optimization for building retrofit: A model using genetic algorithm and artificial neural network and an application. *Energy and Buildings 81*, 444–456.

Costa, A., G. Nannicini, A. Costa, E. Zurich, and G. Nannicini (2018). RBFOpt: an open-source library for black-box optimization with costly function evaluations. *Mathematical Programming Computation* (10), 597–629.

Crawley, B. D. B., C. O. Pedersen, L. K. Lawrie, and F. C. Winkelmann (2000). EnergyPlus : Energy Simulation Program. *ASHRAE Journal* (April), 49–56.

Eisenhower, B., Z. O'Neill, S. Narayanan, V. A. Fonoberov, and I. Mezić (2012, 4). A methodology for meta-model based optimization in building energy models. *Energy and Buildings 47*, 292–301.

Melo, A. P., D. Cóstola, R. Lamberts, and J. L. Hensen (2014). Development of surrogate models using artificial neural network for building shell energy labelling. *Energy Policy 69*, 457–466.

Office for National Statistics (2019). Energy Consumption in the UK (ECUK) 1970 to 2018. Technical Report July 2019.

Prada, A., A. Gasparella, and P. Baggio (2018, 9). On the performance of meta-models in building design optimization. *Applied Energy 225*, 814–826.

Reynolds, J., Y. Rezgui, A. Kwan, and S. Piriou (2018). A zone-level, building energy optimisation combining an artificial neural network, a genetic algorithm, and model predictive control. *Energy 151*, 729–739.

Sousa, G., B. M. Jones, P. A. Mirzaei, and D. Robinson (2017, 9). A review and critique of UK housing stock energy models, modelling approaches and data sources. *Energy and Buildings 151*, 66–80.

Sousa, G., B. M. Jones, P. A. Mirzaei, and D. Robinson (2018, 8). An open-source simulation platform to support the formulation of housing stock decarbonisation strategies. *Energy and Buildings 172*, 459–477.

Wate, P., M. Iglesias, V. Coors, and D. Robinson (2020, 1). Framework for emulation and uncertainty quantification of a stochastic building performance simulator. *Applied Energy 258.*

Wortmann, T. (2018). Genetic evolution vs. function approximation: Benchmarking algorithms for architectural design optimization. *Journal of Computational Design and Engineering 6*, 414–428.