Developing a rapid, scalable method of thermal characterisation for UK dwellings using smart meter data

Jonathan David Chambers

A thesis presented for the degree of

Doctor of Philosophy

June 2017

UCL Energy Institute

Declaration

I, Jonathan David Chambers, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signed:

Abstract

Measuring building thermal performance is important for meeting energy efficiency and emissions reductions goals in the UK. Existing approaches are limited in their capacity to assess the in-situ thermal performance of dwellings, or are expensive and impractically invasive for use on a mass-scale. The UK roll out of smart meters will result in widespread collection of dwelling energy data. However, to-date no established method exists for inferring thermal performance from such data. This research aims to develop a method of measuring the thermal performance of existing dwellings using smart meter data and weather data derived from location which can be applied to large numbers of dwellings.

This research presents the "Deconstruct" method to estimate dwelling thermal characteristics, notably the Heat Loss Coefficient (HLC). This method uses a grey-box building physics model to describe the relation between dwelling thermal properties and measured energy demand. A 'post-hoc control trial' approach identifies data subsets which enable robust model parameter estimates. This enables dwelling characterisation without the need for a controlled experiment with on-site monitoring.

Evaluation against the Energy Demand Research Project (EDRP), Solid Wall Insulation (SWI) field trials, and Pennyland project dwelling datasets demonstrated that the approach and assumptions made in the Deconstruct method were reasonable. HLC values were inferred with a mean uncertainty of 15% for 63% of 7,529 dwellings from the Energy Demand Research Project (EDRP) dataset. The data requirements of Deconstruct were found to be one year of daily average electricity and gas readings, external air temperature, and solar irradiance.

Deconstruct could enable rapid estimation of key thermal performance parameters for millions of UK dwellings, with a range of applications including policy evaluation, retrofit assessments, and advice to dwelling occupants. Future monitoring studies were proposed which address shortcomings in existing datasets.

Acknowledgements

I would like to thank all those who helped and supported me during the course of the PhD.

I would like to thank my supervisors, Prof. Tadj Oreszczyn, Prof. David Shipworth, and Andrew Smith for their guidance and support throughout the PhD journey.

I would also like to extend my gratitude to the academics at the University College London (UCL) Energy Institute for their advice and insights, including Dr. Jenny Love, Dr. Ian Hamilton, Dr. Cliff Elwell, Dr. Alex Summerfield, Simon Elam, and Prof. Bob Lowe. I would also like to thank the excellent UCL Energy Institute admin team for their assistance with many aspects of my PhD.

This research was conducted in part in collaboration with the UK Data Service, who provided access to their Hadoop cluster. I would particularly like to thank Chris Park of the UK Data Service for his technical and moral support on this aspect of the project. I would also like to acknowledge the use of the UCL Legion High Performance Computing Facility (Legion@UCL), and associated support services, in the completion of this work.

This research was made possible by the Research Councils UK (RCUK) Centre for Energy Epidemiology (EPSRC ref EP/K011839/1) jointly funded by the UK Engineering and Physical Sciences Research Council (EPSRC) and Electricité de France (EDF) Energy R&D UK. I would like to thank in particular Dominique Bertin and the R&D team in Hove for their professional support and interest in my work.

This research would not have been possible without extensive input data from a range of providers. These included the UK MetOffice, the US National Centre of Atmospheric Research, NASA, AECOM Building Engineering, the Energy Saving Trust, the UKDA, the Organisation for National Statistics, the Ordinance Survey, DECC/BEIS, the DCLG, EDF, and UCL.

I want to express my appreciation of my family and friends, and most of all Selin for her amazing support, patience, and encouragement over the past years. I would also like to acknowledge the emotional support (however unwitting) of our two cats ‘Evil’ and ‘Confused’.

Acronyms

Name	Description
ADEPT	Annual Delivered Energy, Price, And Temperature
ANOVA	Analysis of variance
ARMA	Autoregressive–moving-average
ARMAX	Autoregressive–moving-average model with exogenous inputs
ASHRAE	American Society of Heating, Refrigeration, and Air Conditioning Engineers
BEIS	Department for Business, Energy & Industrial Strategy
BRE	Building Research Establishment
BREDEM	Building Research Establishment Domestic Energy Model
CARB	Carbon Reduction in Buildings
CDF	Cumulative Distribution Function
CFSR	Climate Forecast System Reanalysis
CHP	Combined Heat and Power
CSV	Comma-separated values
CTSM R	Continuous Time Stochastic Modelling for R
CV	Calorific Value
CVRMSE	Coefficient of Variation of the Root Mean Square Error
DCC	Data Communications Company
DCLG	Department for Communities & Local Government
DECC	Department of Energy & Climate Change
DHI	Diffuse Horizontal Irradiance
DNI	Direct Normal Irradiance
DNI-E	Direct Normal Irradiance-Extraterrestrial
EBC	Energy in Buildings and Communities
EDF	Electricité de France
EDRP	Energy Demand Research Project
EFUS	Energy Follow Up Survey
EHS	English Housing Survey
EPBD	Energy Performance in Buildings Directive
EPC	Energy Performance Certificate
EST	Energy Saving Trust
ETI	Energy Technologies Institute
ETL	Extract, Transform, and Load
EU	European Union
FGHR	Flue Gas Heat Recovery
GHG	Greenhouse Gas
GHI	Global Horizontal Irradiance

Name	Description
GIS	Geographical Information System
GOR	Government Office Region
HDD	Heating-degree-day
HPC	High Performance Computing
IEA	International Energy Agency
IET	Institution of Engineering and Technology
IHE	Integrating the Healthcare Enterprise
IPCC	Intergovernmental Panel on Climate Change
ISO	International Organization for Standardization
JCGM	Joint Committee for Guides in Metrology
KDE	Kernel Density Estimate
LAD	Local Authority Districts
LDZ	Lower Distribution Zone
LIUHB	Loughborough In-Use Heat Balance
MAD	Median Absolute Deviation
MARS	Multivariate Adaptive Regression Splines
MCMC	Markov Chain Monte Carlo
MIDAS	Met Office Integrated Data Archive System
NAC	Normalised Annual Consumption
NAN	Not A Number
NCEP	National Center for Environmental Prediction
NetCDF	Network Common Data Form
NREL	National Renewable Energy Laboratory
NUTS	Nomenclature of Units for Territorial Statistics
NWP	Numerical Weather Prediction
OFGEM	Office of Gas and Electricity Markets
OLS	Ordinary Least Squares
ONS	Office for National Statistics
PDF	Portable Document Format
PRISM	PRInceton Scorekeeping Method
PTG	Power Temperature Gradient
RCUKCEE	Research Councils UK - Centre for Energy Epidemiology
RdSAP	Reduced data SAP
RMSE	Root-Mean-Square Error
SAP	Standard Assessment Procedure
SEDBUK	Seasonal Efficiency of a Domestic Boiler in the UK
SPA	Solar Position Algorithm

Name	Description
SSE	Scottish and Southern Energy
STEP	Seasonal Temperature Energy Price
SWI	Solid Wall Insulation
UCAR	University Corporation for Atmospheric Research
UCL	University College London
UKGBC	UK Green Building Council
UKV	UK variable resolution model
UTC	Coordinated Universal Time
WGS84	World Geodetic System

Nomenclature

Symbol	Description	Unit
A	area	m^2
A_k	area of building element k	m^2
A_{sol}	Effective solar aperture	m^2
C_T	covariance of T_{in} with T_{ex}	-
C_{th}	combined thermal mass of dwelling elements	kJ/°C
C_{ve}	base ventilation coef.	m^3/s
C_{stack}	stack ventilation coef.	$m^3/(s \cdot °C)$
C_{wind}	wind ventilation coef.	m^2
E	energy	MJ
H_{tr}	transmission heat transfer coefficient	W/K
HLC	Heat Loss Coefficient	kW/°C
HTC	Heat Transfer Coefficient	kW/°C
f_s	stack factor	-
f_w	wind factor	-
g	solar energy transmittance of a building element	-
h	height	m
$I_{k,t,tot}$	instantaneous total solar irradiance on element k	kW/m^2
I_{sol}	solar irradiance	kW/m^2
N	number	-
n	exponent	-
P	measured power demand	kW
P_B	measured base-load power demand	kW
P_H	measured energy use for heating	kW
$P_{H,D}$	energy demand for heating	kW
P_{elec}	measured electrical energy demand	kW
P_{gas}	measured gas energy demand	kW
P_{tot}	total measured energy demand	kW
$P_{tot,0}$	Y-intercept of P_{tot} at absolute zero	kW
PTG	Power-Temperature Gradient	kW/°C
Q	quantity of heat	MJ
Q_{in}	heat flow into a zone	MJ
Q_{out}	heat flow out of a zone	MJ
Q_{sky}	net radiative flux between a body and the sky	kW/m^2
q_{stack}	stack volumetric airflow rate	m^3/s
q_v	volumetric airflow rate	m^3/s
q_{ve}	total ventilation airflow rate	m^3/s

Symbol	Description	Unit
T	temperature	$^\circ$C
t	time, period of time	s
T_{ex}	external air temperature (2m from surface)	$^\circ$C
T_{fix}	reference external air temperature for Tin model	$^\circ$C
T_h	base-load to heating regime change point	$^\circ$C
T_{in}	dwelling mean internal temperature	$^\circ$C
T_0	T_{in} at T_{fix}	$^\circ$C
U	thermal transmittance	W/(m$^2\cdot\,^\circ$C)
V	volume of air in a conditioned zone	m^3
v	speed	m/s
α_k	radiant absorptivity	-
β	PRISM house effective heat loss rate	kW/$^\circ$C
ΔP_s	stack pressure difference	mPa
ΔP_w	wind pressure difference	mPa
ϵ	emissivity of a surface	-
η	efficiency, utilisation factor	-
η_B	base-load utilisation factor	-
η_{HS}	heating system efficiency	-
τ	time constant	h
Φ	heat flow rate	kW
Φ_B	heat gain from base-load	kW
Φ_{HS}	heat flow rate from heating system	kW
Φ_i	heat flow rate from source i	kW
$\Phi_{sol,k}$	heat flow rate from solar gains for element k	kW
Φ_{sol}	whole dwelling heat flow rate from solar gains	kW
Φ_{tot}	total heat flow rate	kW
Φ_{tr}	total conductive heat transfer	kW/$^\circ$C
ρ_{air}	air density	kg / m^3
σ	standard error	-
σ_B	Stefan-Boltzmann constant (5.67E-8)	W/(m$^2\cdot$K^4)

Contents

List of Tables

List of Figures

Chapter 1

Introduction

This chapter sets the scene and presents relevant background information against which this study positions itself. The main research question is defined and aims and objectives are laid out. The main structure of the thesis is outlined to provide the reader with guidance for subsequent reading.

1.1 Context

Climate change is the greatest environmental challenge of our time. Limiting greenhouse gas emissions is an urgent problem to keep global average temperature rise below the target of 2 $^{\circ}$ C set by the Intergovernmental Panel on Climate Change (IPCC 2013). Energy use is the largest source of greenhouse gases, making energy savings through efficiency gains one of the most important tools in meeting climate targets (International Energy Agency (IEA 2013)).

The domestic sector in the UK accounts for 25% of energy demand (Department for Business, Energy & Industrial Strategy, (BEIS 2016a)) and 25% of the CO_2 emissions (Cooper and Palmer 2013). The importance of buildings in general, and domestic buildings in particular, in contributing to energy demand has led to European Union (EU) -wide directives aimed at fostering efficiency improvements. These include the 2010 Energy Performance in Buildings Directive (EPBD) (EU 2010) and 2012 Energy Efficiency Directive (EU 2012). These fall under the EU 2020 climate and energy effort (EU 2007) which aims among other targets to improve efficiency across the board by 20% by 2020. The UK has implemented these directives through a range of policies, notably the Climate Change Act of 2008 (HMGovernment 2008), which mandates an 80% reduction in Greenhouse Gas (GHG) emissions by 2050.

Improvements to dwelling energy efficiency are needed to meet environmental challenges and are required under national and international legislation. Space heating in UK dwellings accounts for almost 60% of total dwelling demand (Department for Business Energy & Industrial Strategy

2016), making it a key area for efficiency improvements. However, trends in GHG emissions from buildings suggest that while improvements are being made, the trajectory is not sufficient to meet the mandated 80% reduction (Oreszczyn and Lowe 2010). Although energy use per dwelling per year has fallen from 23,800kWh to 18,600kWh (22%) from 1970 to 2012, total energy use in housing rose by 16% due to the increasing number of homes (Cooper and Palmer 2013). Trends in energy use reduction in homes do not appear on track to meet climate targets in the UK (Oreszczyn and Lowe 2010). Furthermore, flagship policies for improving efficiency such as the Green Deal have not been successful (Rosenow and Eyre 2016).

While a reduction in dwelling energy demand is an aim in itself, it must be remembered that cold, inefficient dwellings have significant impacts on health and also lead to fuel poverty. In England there were an estimated 29,200 excess winter deaths in the 2012-2013, with 21.5% of these (i.e. 6278) being attributed to cold homes, which exacerbate a range of health problems (Wilkinson et al. 2007; IHE 2014). Inefficient dwellings are expensive to heat causing low-income households to find themselves in fuel poverty (IHE 2014). This can place a significant burden on the health and wellbeing of vulnerable households.

The thermal performance of dwellings is of interest to a range of stakeholders including building occupants, building managers (occupants, landlords, councils, etc.), energy utilities, governments and academics. Evaluations of thermal performance enables the condition of existing housing to be assessed and the impacts of energy efficiency policy and retrofit practice to be measured (Dupret et al., n.d.; Roels 2017; Roels 2012), Energy Technologies Institute (ETI 2016). Current approaches to the evaluation of dwelling thermal performance have the following drawbacks.

The Energy Performance Certificate (EPC), which is currently the most widespread assessment of dwelling performance, summarises the expected energy demand, environmental impact, and potential improvements for a dwelling (Figure 1.1). These certificates are a legal requirement for renting and selling dwellings in the UK.

The EPC is generated from a site survey which serves as the input to the Reduced data Standard Assessment Procedure (RdSAP) model, itself a sub-model of the Standard Assessment Procedure (SAP) model (BRE 2011). These models are bottom-up building models and are not designed to deliver as-built performance, although they are frequently (albeit incorrectly) used for this purpose . As a result, there is a gap between the energy demand predicted using SAP/RdSAP and real energy consumption. In addition to the misapplication of SAP for predicting real demand, there are a range of criticisms of the UK EPC, i.e. (Rd)SAP, which may undermine its utility and in some cases its credibility (Scottish Government 2009; Kelly, Crawford-Brown, and Pollitt 2012; Hughes et al. 2013; Stone et al. 2013; Li et al. 2014; Dickson et al. 1996; Doran and Carr 2008; Shipworth 2011; Cooper and Palmer 2013; Department of Energy and Climate Change 2014; Summerfield, Oreszczyn, et al. 2015a). The onsite inspection required to deliver an EPC is intrusive which, together with costs between £60 and £120 (USwitch 2017), precludes regular

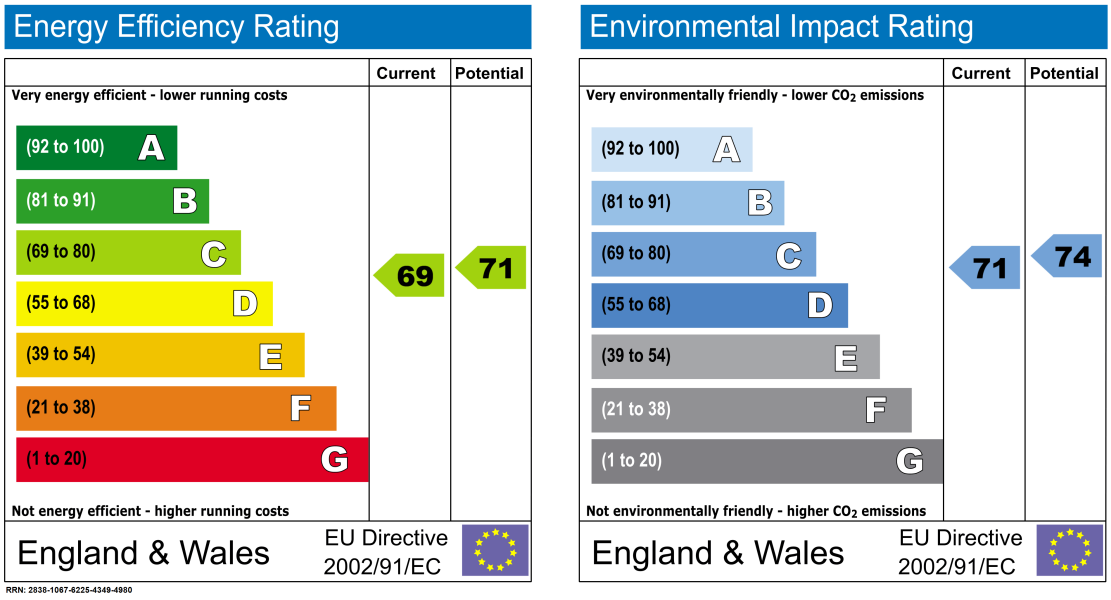

Figure 1.1: Example of energy and emissions ratings from an EPC report

large scale evaluation of dwelling condition.

The co-heating test is currently the empirical method used for measuring whole-house fabric thermal performance through in-situ experimentation on a given dwelling; it is used to measure the Heat Transfer Coefficient (HTC) (kW/$^\circ$C) of a house (Everett 1985; Stamp 2015). Given a good experimental protocol, this approach results in estimates of dwelling thermal performance with an uncertainty of $\pm 10\%$ (Jack 2015; Stamp 2015). However, this method requires the entire building to be vacated for a period of at least two weeks during the winter and requires significant equipment and skill for it to be performed correctly. It is therefore by definition limited in its application.

Direct measurement of the thermal properties of individual elements of a dwelling can also be performed, for example using heat flux sensors to measure wall U-values (Biddulph et al. 2014). These however can only assess single dwelling elements and do not necessarily give a good overall picture. Similarly to co-heating, they are expensive and intrusive to perform.

In the face of these challenges, new developments in energy metering technology and data collection, notably the installation of smart meters in UK dwellings, enable new approaches to measuring energy demand (DECC 2012).

Smart meters are advanced energy meters that are able to record and store consumption data in real-time, and transmit this to remote servers enabling rapid retrieval of dwelling data. The remote connection is in some cases bidirectional, allowing the meter to potentially receive messages and initiate actions based on their contents (Depuru, Wang, and Devabhaktuni 2011). These meters are being installed across Europe including the UK, where the deployment is expected to cover

most dwellings by 2020 (BEIS 2016b). Globally, smart meter installations have been expanding rapidly with a global penetration expected to reach 53% by 2025 (Navigant Research 2016).

Smart meters are intended to provide a number of benefits derived from enabling the remote collection of electricity and gas consumption data. They can enable improved billing based on true consumption instead of estimates, as well as help reduce fraud and disputes (IET 2013; Ofgem 2013). They could support time-of-use tariffs for demand shifting (IET 2013) as well as motivate behaviour change through other forms of feedback (Beard 2015; Ofgem 2013; Raw, Ross, and Cripps 2011; Kim and Shcherbakova 2011). In the context of this research they provide the raw input for data-driven energy demand models, with no extra cost for data collection eliminating the need for bespoke monitoring campaigns.

1.2 Problem statement

The challenges of obtaining reliable performance estimates of dwellings has led to widespread observations of a 'performance gap' between expected and measured total energy demand. This has been taken by some as evidence for the rebound effect in dwelling retrofits whereby cost savings from efficiency improvements lead to increased demand for energy services (notably higher indoor temperatures) (Greening, Greene, and Difiglio 2000), but is likely better described as a 'credibility gap' (Bordass, Cohen, and Field 2004; D'Oca and Hong 2014) caused by limitations of the assessment methods.

There is currently no established method for estimating the as-built thermal performance of dwellings in the UK in a manner which can be cost-effectively scaled to large numbers of dwellings. Some research has made progress towards in-situ measurement of dwelling thermal performance including work in Meyer (1981); Fels (1986); Rabl (1988); ASHRAE Handbook (2005); Bacher and Madsen (2013); Summerfield, Oreszczyn, et al. (2015b), Jack (2015) which this research aims to build on.

1.3 Research Question

In order to address the challenges described in Sections 1.1 and 1.2 raised by the literature review, this research sets out to answer the following research question:

Q How can "as-built" thermal characterisation of UK dwellings be performed rapidly and non-intrusively at scale?

"Thermal characterisation" is defined in this thesis as the distinctive features of dwelling thermal demand, such as thermal loss and gain coefficients, and the values of parameters describing those

features. By “rapid” is meant that parameter inference should not require significant man-hours or computing time per dwelling. By “at scale” is meant the ability to process a large and growing number of dwellings. This drives the need for non-intrusive methods to avoid disruptive and costly monitoring campaigns.

1.4 Aim and Objectives

The overall aim of this research has been to develop and evaluate a method of measuring the thermal performance of existing dwellings. This method should be scalable such that it could be applied to the large repositories of smart meter data collected by energy utilities and generate performance profiles for large numbers of these dwellings with minimal manual intervention. The aim is for the method to scale such that increasing the number of dwellings studied does not result in a significant increase in the required man-hours or computing time. However, the aim is not to minimise the monitoring time, relying instead on the ease of archival of long smart meter data time series.

This research is concerned with developing a suitable method, assessing its uncertainty, and comparing it to existing methods. It aims to establish the parameters necessary to characterise thermal performance, as well as how to infer parameters from data, and what the requirements on input data are.

A further objective is to establish that the method is able to estimate thermal performance parameters of significant numbers of dwelling, and to contrast these results with outcomes of other methods notably EPCs. This should also determine conditions which preclude a dwelling from being reliably assessed by this method.

1.5 Scope

The thermal performance of a dwelling is a concept that can be understood in a variety of ways. For some, it is thermal properties of the physical components only, while for others it is the ability of the dwelling to provide thermal comfort. This research tends towards the former definition rather than the latter, as thermal comfort of a person is itself a complex subject (Fanger 1970). For simplicity, a concept of thermal performance was chosen along the lines of that used in co-heating tests - the thermal performance of a dwelling is its ability to retain thermal energy within its boundary. However, unlike co-heating tests the efficiency of the heating system must be taken into account.

The thermal performance of dwellings is defined in terms of the thermal processes that govern heat transfer in and out of the dwelling envelope. These can in turn be measured (operationalised)

through the relation between metered energy input to the dwelling and the weather conditions which drive the thermal processes.

This research is limited to dwellings in the UK for which both electricity and gas data is available. These dwellings are drawn from several datasets of dwelling energy demand. Where such metadata is available, dwellings which are known to use heating technologies such as Combined Heat and Power (CHP) units or heat pumps are excluded. Energy demand is considered on a per-dwelling basis, as opposed to per occupant or unit area. A dwelling unit is defined as being the entity monitored by a single meter point.

This research focuses on the physical properties of buildings. It aims to determine physical heat loss and gain coefficients independent of occupant effects.

For the purpose of this research, smart meters are defined as dwelling energy meters that can record demand at least at hourly frequency, store that information, and transmit it to a remote data repository. Although certain definitions of smart meters include requirements for being able to send feedback information to the meter/dwelling, this aspect of smart metering is not relevant here.

This research uses weather data, including but not limited to temperature, solar irradiance, and wind speed, which is selected to match the location and time span of a given dwelling's energy data. Weather data is used from back-corrected numerical weather prediction models, as being the closest approximation to local weather conditions in the absence of on-site weather monitoring.

1.6 Contribution

The principal contribution of this thesis will be a novel method for estimating dwelling thermal performance. The method will enable the study of individual dwellings, notably for measuring the impact of retrofits and examining the performance/credibility gap. The method will also enable large numbers of dwellings to be assessed at low marginal cost, as it makes opportunistic use of existing investment in smart metering infrastructure as well as using widely available weather data.

1.7 Outputs

The outputs of this thesis are listed in this section.

1.7.1 Papers

J. Chambers, T. Oreszczyn, and D. Shipworth, "Quantifying Uncertainty In Grey-box Building Models Arising From Smart," Build. Simul. Conf., vol. 0, no. 1, pp. 2947–2954, 2015. (Chambers, Oreszczyn, and Shipworth 2015)

J. Chambers, V. Gori, P. Biddulph, I. Hamilton, T. Oreszczyn, and C. Elwell, "How solid is our knowledge of solid walls? - Comparing energy savings through three different methods," in CISBAT 2015 - International Conference Future Buildings & Districts Sustainability from Nano to Urban Scale, 2015. (Chambers et al. 2015)

1.7.2 Computer programs

J. Chambers and A. Stone, "EPC toolkit." RCUK-CEE, 2016. https://github.com/RCUK-CEE/epctk (Chambers and Stone 2016)

1.7.3 Datasets

This thesis has produced transformed versions of primary dwelling energy datasets which have been subjected to extensive cleaning, standardisation, and addition of relevant metadata. To the extent permissible under data protection agreements, these datasets will be made available to a wider audience.

1.8 Chapter outlines

1.8.1 Chapter 2. Literature review

This chapter presents the environmental and policy context of thermal efficiency in dwellings, as well as the current state of the art of modelling existing dwelling thermal performance. An overview is given of the state and potential of smart meter infrastructure. Next, a review is made of existing methods for whole house thermal assessment.

1.8.2 Chapter 3. Theory

This chapter outlines the physical theory underpinning the main thermal transfer processes in dwellings with focus on empirical approximations for whole-dwelling thermal flows. These include conductive and radiative heat exchange, heat loss through ventilation, as well as internal gains. A whole-dwelling thermal energy model is given. A formulation of the Heat Loss Coefficient in terms

of measured energy demand is given taking into account linking functions between thermal flows and metered energy use.

1.8.3 Chapter 4. Method

This chapter outlines the method developed to enable reliable estimation of model parameters from metered data and weather data. An in-depth explanation of the data treatment required for this is given. An approach for estimating the model coefficients for an occupied dwelling is presented.

1.8.4 Chapter 5. Datasets

This chapter details the datasets used for the development and evaluation of the method developed in this thesis, and notably the extensive data processing in order to transform raw input files into useable data. This covers the smart meter data, weather data, and dwelling survey data used. Quantitative descriptions and comparisons of the datasets used are given.

1.8.5 Chapter 6. Method Evaluation

This chapter presents the evaluation of the precision, accuracy and reliability of the method. Evaluations are made of key modelling assumptions. Model uncertainties are calculated using a Monte-Carlo approach. The ability of the model to predict year-on-year consumption is assessed.

1.8.6 Chapter 7. Results

This chapter presents several case studies making use of the method developed in this thesis. The first case study was a pilot project performed in collaboration with the UK Data Service using the EDRP smart meter dataset to estimate the thermal performance of a large number of dwellings using a Big Data cloud computing system. The second case study investigates dwelling type classification using inferred thermal properties using data from the Pennyland project. The third case study aims to explore the dwelling characteristics (type, occupancy, ownership, etc) as determinants of thermal performance, and to classify dwellings by those characteristics.

1.8.7 Chapter 8. Discussion

This chapter discusses the current research findings with respect to previous research and describes potential implications for whole dwelling thermal assessment in terms of policy and practice in thermal efficiency of the UK residential sector.

1.8.8 Chapter 9. Conclusions

This chapter presents a summary of key findings from the research and a discussion of the contributions to knowledge. Limitations of the current research are highlighted and potential areas of future research are identified.

Chapter 2

Literature Review

2.1 Introduction

This chapter first presents the environmental and policy context of thermal efficiency in dwellings. The need for dwelling thermal performance evaluation and current challenges to this evaluation are discussed. Then, a review of the drivers of smart metering infrastructure rollout and the associated opportunities is given. It also reviews whole house thermal performance assessment methods , covering government-standard EPC assessments, co heating tests, in-situ monitoring, infra-red thermography, and methods for assessing performance from meter and sensor data. This chapter also provides relevant background to the building physics theory, presented in Chapter 3.

2.2 Environmental and policy context

Limiting greenhouse gas emissions is an urgent problem to keep global average temperature rise below the target of 2 $^{\circ}$ C (IPCC 2013). Worldwide, the built environment accounts for 40-50% of natural resource use, 20% of water use, 30-40% of energy use and around a 33% of CO_2 emissions (Cooper and Palmer 2013). Energy efficiency improvements have been identified as a major potential contributor to emissions reductions with up to 48% of total CO_2 abatement coming from efficiency improvements by 2035 in the IEA's 450ppm scenario (Figure 2.1 IEA (2013)).

The UK dwelling stock consists of approximately 27 million dwellings and is responsible for 26.6% of final energy consumption (Figure 2.2) and 27% of CO_2 emissions (Cooper and Palmer 2013). Space heating is estimated to account for around 57% of housing energy use (BEIS 2016a).

Inefficient dwellings also have significant impacts on public health and wellbeing in the UK. In England there were an estimated 29,200 excess winter deaths in the 2012-2013 winter with approx-

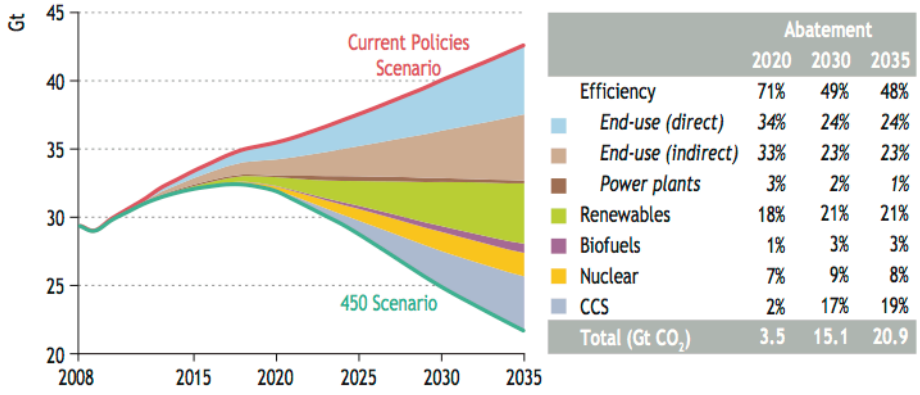

	Abatement		
	2020	2030	2035
Efficiency	71%	49%	48%
End-use (direct)	*34%*	*24%*	*24%*
End-use (indirect)	*33%*	*23%*	*23%*
Power plants	*3%*	*2%*	*1%*
Renewables	18%	21%	21%
Biofuels	1%	3%	3%
Nuclear	7%	9%	8%
CCS	2%	17%	19%
Total (Gt CO_2)	3.5	15.1	20.9

Figure 2.1: Contributions to total CO_2 emissions reductions from various sectors according to IEA World Energy Outlook 450ppm Scenario IEA (2010) p.394.

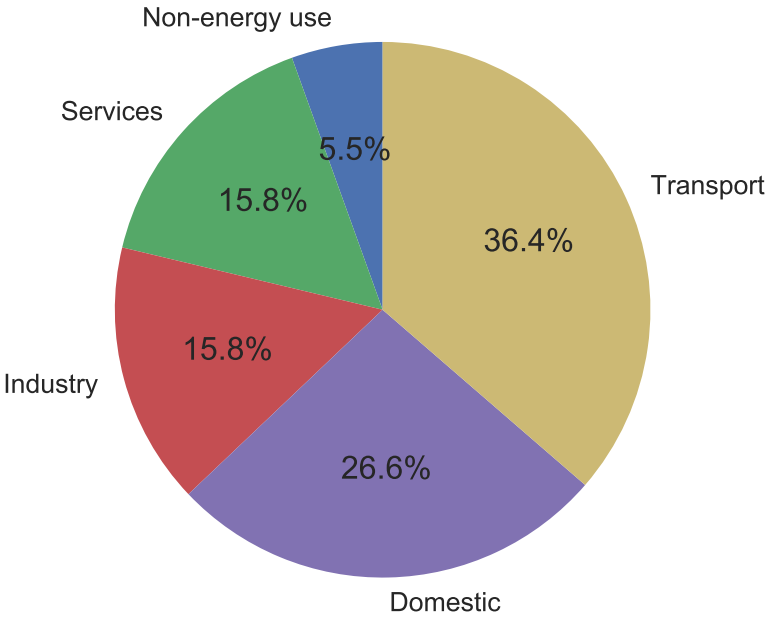

Figure 2.2: UK final energy consumption by sector 2016 BEIS (2016a)

imately 21.5% (6278 deaths) attributed to the coldest homes (IHE 2014). Cold homes exacerbate a range of health problems including respiratory and circulatory conditions (Wilkinson et al. 2007). Most fuel-poor households occupy low quality dwellings and are exposed to cold, humidity, and mould, with negative health impacts (Oreszczyn et al. 2006) - efficiency, poverty, and health were found to be closely interlinked (Gilbertson and Green 2008). Efficiency improvements in this sector are therefore critical to reducing total consumption and pollutant emissions, and improving quality of life (Cooper and Palmer 2013).

The European Union (EU) agreed in 2007 to improve energy efficiency by 20% by 2020 (EU 2007), which was later extended to 40% reductions in greenhouse gases and 27% efficiency improvement by 2030 (EU 2014). This created a need for suitable compliance processes and energy rating tools, for which the EU has put in place several mechanisms (Berry et al. 2014). The Energy Performance in Buildings Directive (EPBD) focuses on building energy demand. A headline goal of the EPBD is that from 2020 new buildings in the EU should consume nearly zero energy. Member states should encourage retrofits, and minimum standards for various building components were introduced. Importantly, efforts have been put into harmonising efficiency calculation methods. Among other measures, member states are required to develop a method of measuring and displaying building energy consumption through the Energy Performance Certificate (EPC), with the long-term goal of achieving net-zero-energy buildings (EU 2010).

In addition to its European commitments, the UK has enshrined an 80% reduction in emissions by 2050 in the 2008 Climate Change Act (Parliament 2008). Further proposals include an 80% reduction specifically in UK housing emissions by 2050, as well as for all new homes to be 'zero carbon' by 2016, though these targets have been modified to target non-domestic by 2019 while private dwellings are still under discussion (Boardman, Bank, and Britain (2007) and UK Green Building Council (UK-GBC (2014))). A similar proposal to cut greenhouse gas emissions in existing homes by 29% by 2020 through whole-house retrofit was put forward by the Department of Energy & Climate Change (DECC) in 2010 (DECC 2009).

To date the UK policies have increased the percentage of dwellings with insulation with on average 2.5 efficiency measures of some kind per dwelling, but equally left many dwellings un- or under-insulated (Hamilton et al. 2014; DCLG 2013; Stafford, Gorse, and Shao 2011). Each dwelling consumed on average 4226 kWh of electricity and 15257 kWh of gas in 2012 representing a 21% reduction per dwelling compared to 1970 (DECC 2013a). The Annual Delivered Energy, Price, and Temperature (ADEPT) model and Seasonal Temperature Energy Price (STEP) model (Summerfield, Lowe, and Oreszczyn 2010) identify the trajectory of total delivered energy to UK households. These found that changes in energy demand were consistent with changes in temperature and energy price alone - there appeared to be no trend of increasing energy efficiency beyond the (limited) price elasticity of consumers. Improvements in energy efficiency in dwellings are not distinguishable from the effects of warmer winter temperatures and occupant responses to energy

costs. Considerable challenges therefore remain to drive a reduction in energy demand from homes.

2.3 State of dwelling thermal performance assessment

Measuring the thermal performance of existing dwellings is important in several contexts related to meeting efficiency improvement goals including:

- evaluating the condition of existing housing (Dupret et al., n.d.; Roels 2017)
- measuring the impacts of retrofits at the level of individual buildings as well as building stock-level changes due to retrofit campaigns (Roels 2012; ETI 2016)
- implementing performance-based contracting for the delivery of energy efficiency services
- evaluating the impact of policy and the evolution of socioeconomic context on thermal performance trends.

Evidence suggests that predictions of performance, based upon a summation of expected elemental performance, are prone to significant inaccuracy – the so-called 'performance gap' (Jack 2015). Measured performance is mostly found to be worse than expected, although there appears to be a trend where the most inefficient dwellings perform somewhat better than expected (Sunikka-Blank and Galvin 2012; Allibe 2012).

'Bottom-up' models of energy demand are widely used for performance assessment, for example the SAP model which is the basis of EPC assessment in the UK (BRE 2011; Jack 2015). These depend on accurate estimates of the thermal properties of dwelling components such as walls - a number of studies have found significant differences between the expected and measured performance of these components. Doran and Carr (2008) measured cavity wall U-values before and after cavity wall insulation installation using heat flux sensors for about 100 dwellings. They found that the improvement in U-value was on average 38% less than the expected improvement calculated using the theoretical properties of the cavity wall and insulation material. In an experimental study, Li et al. (2014) determined the wall U-value of 58 solid-walled dwellings (40 brick solid walls and 18 stone walls, generally considered to be the least efficient construction type) using a lumped thermal mass and inverse parameter estimation technique. They found the mean U-values to be $1.3 \pm 0.4 Wm^{-2} K^{-1}$ compared to the standard RdSAP assumption of $2.1 \pm 0.4 Wm^{-2} K^{-1}$ - a 38% reduction in the expected heat flux rate and correspondingly a better than expected performance for solid walls. These results indicate that there can be significant differences between expected and actual wall U-values, which in turn have a significant influence on total building heat loss rates.

Similarly, whole-dwelling assessments have found significant differences between expected and measured consumption. Stafford, Bell, and Gorse (2012) reported that in 30 out of 34 whole-dwelling assessments using co-heating the measured dwelling heat loss was greater than predicted,

with discrepancies ranging from 1% to over 120%.

Several studies describe this gap as a 'credibility gap' (Bordass, Cohen, and Field 2004; D'Oca et al. 2014), caused by assumptions made in assessments not being sufficiently representative of reality. This suggests an alternative interpretation of observations than the 'performance gap' terminology. Using the term 'performance gap' may be misleading in implying that the dwelling construction is to blame for the discrepancy between measured and predicted consumption, when in fact limitations in the prediction methods may be the cause. Such discrepancies have also formed much of the basis for the discussion of rebound effects in thermal retrofits (Sanders and Phillipson 2006; Rosenow and Galvin 2013). The principle of direct rebound is that improving efficiency reduces energy costs, leading people to use more energy now they can afford it. It has been sometimes assumed (Deurinck, Saelens, and Roels 2012) that failure to achieve energy reduction goals has been the result of occupant behaviour, through increased demand temperatures. However evidence for occupant-driven post-retrofit temperature change beyond the expected physical changes is weak (Deurinck, Saelens, and Roels 2012). Instead, inadequacies in the most common performance assessment methods may be the source of the observed mismatch.

There is therefore a need to develop improved methods of assessing the thermal performance of existing dwellings, which motivated the development of Energy in Buildings and Communities (EBC) Annex 71 'Building Energy Performance Assessment Based on In-situ Measurements' (Roels 2017). This work is facilitated by the emergence of improved energy data collection through smart metering initiatives.

2.4 Smart Meters

The European Commission defines an intelligent metering system or 'smart meter' as an electronic device that can measure the consumption of energy, and can transmit data using a form of electronic communication. A feature of a smart meter is the ability to provide bi-directional communication between the consumer and supplier/operator - however from the perspective of data-driven analytics it is the ability of the meter to transmit consumption data directly to an external data repository that is most important (EC 2010).

Smart meters are being rolled out across the EU in accordance with the 2006/32/EC Directive (EU 2006). This Directive supports the delivery of energy-saving programmes and other measures aimed at improving end-use energy efficiency. The EU aims to achieve 70% penetration of smart meters across member states by 2020 (EC 2014). The UK aims for nearly all 27 million homes to have smart meters by 2020, as of September 2016 there were 4.05 million smart meters (1.69 million gas and 2.35 million electricity) operated by large energy suppliers in domestic properties across Great Britain (Office for National Statistics (Experimental Office of National Statistics 2016)). Smart meter data in the UK is channelled through the Data Communications Company (DCC),

a regulated monopoly that acts as a gateway between consumers' meter data and relevant users of that data - including the user's energy supplier, the network operators, and third parties which may be authorised by the consumer to access their data - for example companies offering energy related services through smart appliances (Figure 2.3) (Beard 2015, DCC (2017)).

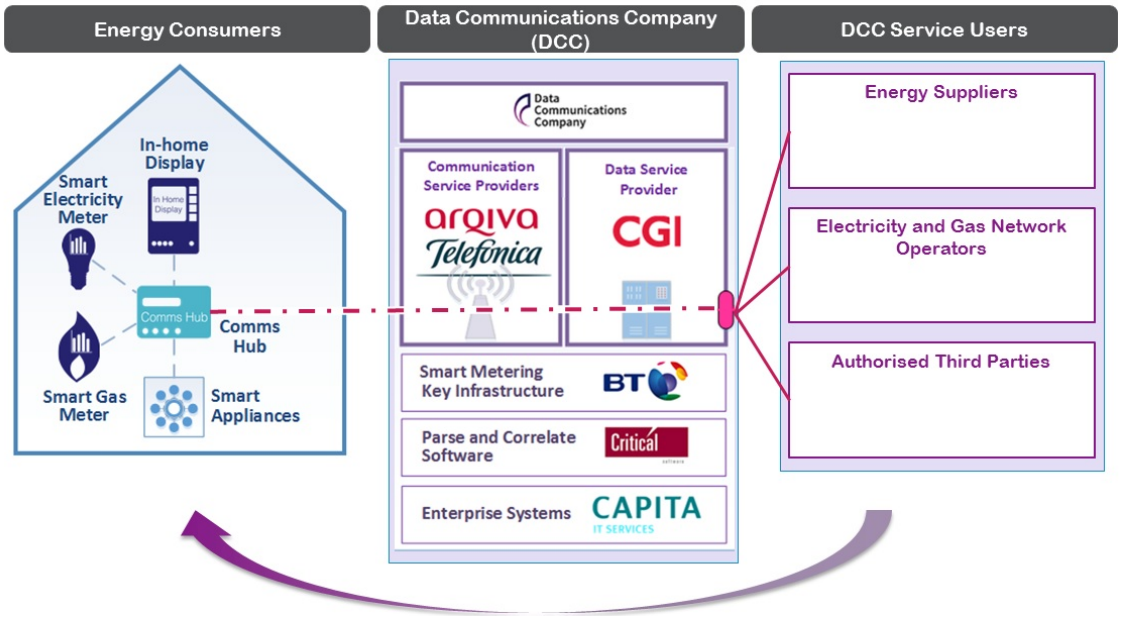

Figure 2.3: Overview of the DCC infrastructure connecting domestic smart meters to energy suppliers, network operators, and third parties (Data Communications Company 2017)

The technical capabilities of UK smart meters are defined by the Government in DECC (2011). The smart meter devices installed on electricity and gas lines are capable of recording energy in kWh with a date-time stamp at half hourly intervals. The devices also hold billing data (tariffs and prepayment information). In addition, the meters are capable of calculating and storing energy consumption:

- daily total for the current day and previous eight days
- weekly total for the current week and the preceding five weeks
- monthly total for the current month and the preceding thirteen months

Currently only monthly information is sent by default to the energy supplier for billing purposes only. Access to more detailed data requires a signed consent form from the customer (Beard 2015).

Smart meters are intended to provider a number of benefits:

- Smart meters enable the collection of electricity and gas consumption data, remotely in real-time and without manual intervention
- Clear and accurate billing, including the prevention of fraud and the reduction in consumer-utility disputes (IET 2013; Ofgem 2013; Moon, Hollings, and Nop 2014)
- Support time-shifting of energy use, for example through time-of-use pricing (IET 2013)

- Motivate behaviour change in consumers toward more efficient behaviours by providing real-time information about energy use (such as energy displays) and enabling services based on this information (Beard 2015; Ofgem 2013; Raw, Ross, and Cripps 2011; Kim and Shcherbakova 2011)
- Use in research context, for example as input to data-driven energy demand models (Albert and Rajagopal 2013), with no extra cost for data collection

Smart meters present an opportunity for improved dwelling assessment as they provide up-to-date, accurate energy demand information for large numbers of dwellings. The dwelling metered energy demand can be used for a range of modelling approaches. As government and energy utilities have committed funds and infrastructure to the installation of smart meters, analysis can be performed at low marginal cost as there is no additional cost for data collection above that already in place. In fact, energy utilities are actively seeking new ways to draw value from the data collected (Laurent 2014).

2.5 Thermal performance assessment methods

This section reviews methods for determining whole dwelling thermal performance based on dwelling designs, dwelling on-site data collection (surveys, monitoring using sensors), as well as using measured energy use and thermal flows.

2.5.1 The EPC and SAP model

The Energy Performance Certificate (EPC) is currently the most widespread assessment of dwelling performance, as it is integral to building efficiency legislation under the EPBD (EU 2010) and a legal requirement for renting and selling dwellings in the UK. The EPC summarises the expected energy demand, environmental impact, and potential improvements for a dwelling.

The Standard Assessment Procedure (SAP) model is the UK standard dwelling energy model, based on the Building Research Establishment's (BRE) Building Research Establishment Domestic Energy Model (BREDEM) (BRE 1997). It is also closely related to International Organization for Standardization (ISO)-13790, which describes a standard for the calculation of space heating and cooling of buildings and which is the basis for or related to a number of dwelling assessment methods in the EU (ISO 2010). SAP describes the inputs, calculations, and assumptions needed to estimate the energy demand of a dwelling based on its design or surveyed building properties (e.g. wall types and composition, windows), equipment, and occupancy, and can be applied to a building at design stage or an existing building when all design details are known (BRE 2011). The underlying BREDEM model uses a mixture of analytical techniques (calculating energy balances) and empirical functions to predict energy consumption for water heating, cooking, lights and

appliances (Anderson et al. 1997), an illustration of the components of the energy balance is shown in Figure 2.4.

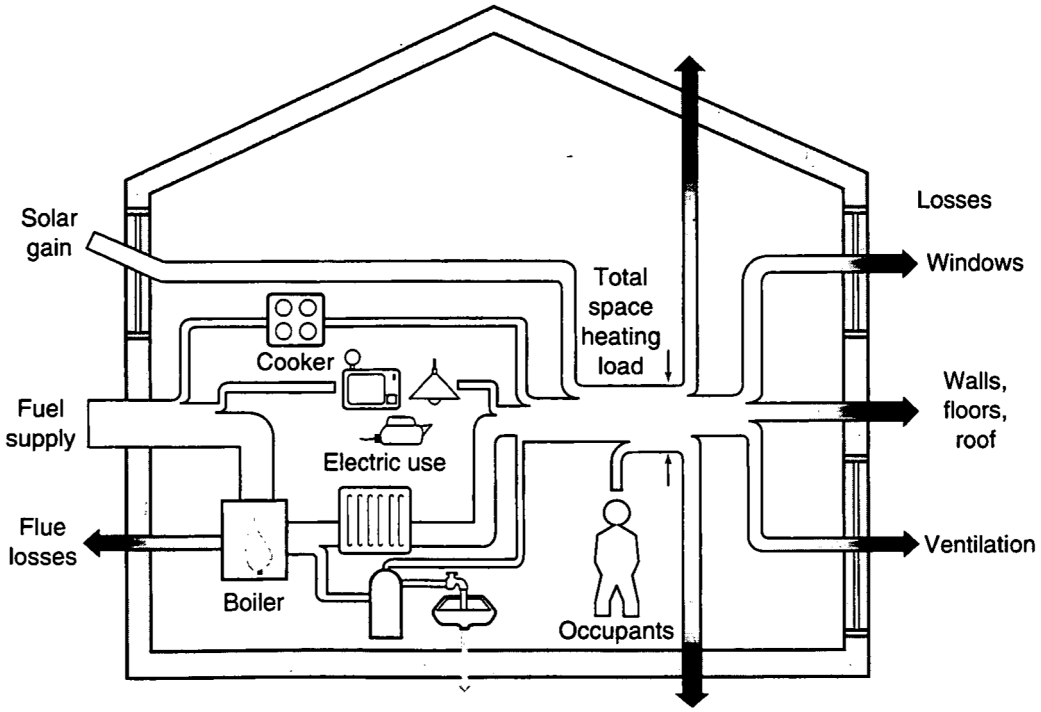

Figure 2.4: Schematic illustration of the energy balance principle used in BREDEM-8 p.5.

The Reduced SAP (RdSAP) is used for the EPC and describes a set of additional assumptions and calculations to generate inputs for the SAP model for an existing dwelling for which all design details are not known - for example deriving wall U-values from building age and observed construction type (BRE 2011).

The EPC delivers the following information to the end-user (UK 2012):

- A SAP rating on a scale of 1 - 100, based on the energy costs (monetary costs) associated with space and water heating (taking into account savings from energy generation), adjusted for floor area.
- An efficiency class from A (best) to G (worst) corresponding to SAP rating bands.
- An Environmental Impact rating (EI) on a scale of 1 - 100 based on the annual CO_2 emissions, adjusted for floor area. This takes into consideration the emission intensity of the various fuels required.
- A dwelling CO_2 emission rating in kg m^{-2} $year^{-1}$, equal to the annual CO_2 emissions per unit floor area (taking into account savings from energy generation).
- Additional advice regarding what retrofits might be cost effectively implemented (Which 2014).

These energy ratings depend heavily on the thermal performance of the dwelling through the heating energy demand, but also take into account energy used for hot water, heating, lighting, and appliances. As such, it is important to keep in mind that the EPC is not exclusively an assessment of dwelling thermal performance.

The RdSAP based EPC, as well as regular SAP models are not designed to deliver in-use performance, although they are frequently (albeit incorrectly) used for this purpose (Oreszczyn 2014). Criticisms of EPC and (Rd)SAP have been given including:

- SAP uses outdated assumptions for prevalence and utilisation of various technologies (Scottish Government 2009).
- SAP assumes dwellings are adequately heated and that heating systems are adequately sized and assumes a standard two-period heating pattern, while in a study using Carbon Reduction in Buildings (CARB) project data it was found that only 40% of dwellings displayed a heating pattern corresponding to the 2-peak (morning and evening) profile assumed (Huebner et al. 2014).
- SAP makes assumptions about demand temperatures that are not consistent with observations. Analysis of CARB and EFUS datasets found that only 20% of monitored dwellings ever reached the assumed demand temperature of 21 $^\circ$C, while the mean demand temperature was 19.3 $^\circ$C (Huebner et al. 2013; DECC 2013b). There was also much less difference between weekday and weekend heating patterns than expected in SAP (Huebner et al. 2013). SAP energy rating is particularly sensitive to the demand temperature (Hughes et al. 2013; Hughes et al. 2014).
- Some of the parameter assumptions made in RdSAP are incorrect, such as U-values for solid walls (Li et al. 2014). The SAP model output is particularly sensitive to the U-value, as well as other parameters for which potentially incorrect assumptions are made (Hughes et al. 2013).
- The EPC embodies multiple, sometimes contradictory, policy aims, notably conflating lowering energy cost with improving efficiency, which has made unclear the real meaning of reported energy demand ratings (Kelly, Crawford-Brown, and Pollitt 2012).
- Current legislation only requires EPC certificates to be created or updated when the property is built or put up for rent or sale (UKGovernment 2017), resulting in EPCs potentially being outdated. Notably, there is no requirement for a new EPC to be delivered after dwelling alterations (energy-related or otherwise).
- EPCs are created manually by an inspector which introduces human errors and results in a lack of reproducibility of results. A 'mystery shopper' experiment which invited multiple evaluations of the same dwellings found significant variation in the EPC ratings produced by the different assessments. The range of EPC ratings spanned at least two EPC bands for almost two thirds of the dwellings analysed (Department of Energy and Climate Change 2014).

- Manually creating an EPC for every dwelling is costly, time consuming, and inconvenient.

The weight of evidence suggests that the EPC is not fit for the purpose of providing an accurate indication of the performance of an existing dwelling (Kelly, Crawford-Brown, and Pollitt 2012). This should not be surprising as SAP by design aim to provide normalised calculations using common assumptions. Nevertheless its limited ability to capture the real performance of a dwelling limits its utility in addressing the challenges in dwelling thermal performance and have helped undermine retrofit policies such as the failed UK Green Deal (Rosenow and Eyre 2016).

Similar calculation approaches to RdSAP are used to generate EPCs in other EU countries, including the Netherlands. A study of large databases of EPCs compared to measured annual consumption demonstrated a significant trend whereby older dwellings with low efficiency ratings tended to consume less energy than expected, while new dwellings with high efficiency ratings consumed more than expected (Majcen, Itard, and Visscher 2013). This effect has been labelled as the 'prebound' effect, has as also been observed in other studies on Swiss, German, Dutch, French, and Belgian dwellings, all of which use bottom-up calculation methods related to ISO-13790 (Sunikka-Blank and Galvin 2012, Hoffmann and Geissler (2017)). The prevalence of this issue indicates a basic problem with the assessment approach, which is significant because as a rule retrofits and policy interventions are based on EPC labels rather than measured energy data.

It is important to highlight again that in the UK and other aforementioned countries the building assessment is performed manually. As a result, the quality and accuracy of the assessment will be ultimately limited by the skill and attention to detail of the assessor. In the UK the requirements to become an assessor are relatively light. A certification may be obtained through distance learning or a 3-5 day course with no requirement for prior experience qualifications, and EPCs delivered by a certified assessor are not usually cross checked (EST 2017). It is therefore reasonable to suspect that may assessments may be of low quality. This is clearly the case in the aforementioned 'mystery shopper' experiment - if the results for a given dwelling range over two efficiency bands, they cannot all be correct.

This human aspect ultimately limits the impact of possible improvements to SAP and RdSAP. SAP has undergone a number of revisions since the original 1995 version, including revisions in 1998, 2001, 2005, 2009 and 2012, which have introduced new elements such as monthly energy calculations and region-specific climate profiles (BRE 2014). These have also been reflected in RdSAP. However while these changes can improve the theoretical and physical accuracy of the building description, they cannot solve the problem of human error.

2.5.2 In-situ U-values

The U-values of building elements can be measured in-situ using heat flux sensors affixed to building elements, with experimental protocol recommendations being given by ISO 9869:2014. This

experimental approach has been applied to test-cells for monitoring specific building components as well as to existing dwellings (Doran and Carr 2008; Baker and Dijk 2008; Birchall, Pearson, and Brown 2011). Using steady-state analysis, typically over 7 days of monitoring are required, with measurement durations requiring 24 hour integer periods, across similar periods of cold weather. Heat flux measurements can also be made in occupied dwellings with varying heating patterns, as was done in the Solid Wall Insulation trials (Birchall, Pearson, and Brown 2011). Alternatively, a number of dynamic analysis protocols can be used (Biddulph et al. 2014; Bacher and Madsen 2011). Dynamic approaches can obtain results much more quickly as well as provide information on the dynamic properties of the built element such as heat capacitance. These approaches require good quality time series data and as such tend to require extensive manual pre-processing (Gori 2016; Madsen 2015).

While element U-value measurements may be highly accurate, they are not sufficient for an overall assessment of dwelling thermal performance. Measurements are localised to a small section of the dwelling and cannot take into account doors, windows, or any kind of ventilation driven losses.

Chambers et al. (2015) compared a SAP annualised energy demand calculation with assumed wall U-values and also U-values measured by heat-flux sensor against weather-normalised metered total energy demand. It was found that the annual consumption estimates were similar for measured and assumed wall U-values, even though the measured U-values sometimes differed significantly from assumptions. By contrast there was little correlation between metered energy demand and SAP demand, suggesting that the U-value alone is not a good indicator of overall performance.

2.5.3 Infra-red thermography

Infrared thermography allows visual checks on heat loss across the building fabric. The use of thermography is increasing as prices reduce and equipment becomes more user friendly. The method is governed by international standards (ISO 13187:1999) (Stamp 2015). Thermography allows the rapid identification of defects, with the proviso that conditions at the time of image capture are appropriate and operatives with suitable experience can analyse the results, which may otherwise be misleading. Thermography can be used as a bulk technique, for example by making aerial infra-red scans of urban areas (Gregg 2014), or for performing street-level imaging such as in the Bristol CHEESE project (Andrews 2017). Improvements in thermal imaging technology have facilitated techniques such as fly-by (aircraft-mounted) and pass-by (car mounted) imaging. These are useful as a qualitative indicator of relative performance and may help identify the presence of defects (Fox et al. 2014). However, the resulting images may serve better as motivational tools to encourage retrofits than as quantitive measurements of energy loss rates, since a range of factors make it difficult in practice to link thermal images to dwelling heat loss coefficients (Jack 2015).

2.5.4 Co-heating

The co-heating test is a method for assessing whole building fabric performance - that is, the aggregate performance of building components. The principle of the co-heating test is based upon an assumed steady-state energy balance; its purpose is to measure the total (both fabric and infiltration) heat loss rate of a building, with the result most commonly reported as a heat loss or a Heat Transfer Coefficient (HTC) with units of Watts per Kelvin and is based on the method developed by Siviour (1981) and Everett (1985).

During a co-heating test, the investigated dwelling is homogeneously heated to an elevated steady-state interior temperature (usually 25 $^{\circ}$C) using electric heaters and fans. The electrical energy use, indoor and outdoor air temperatures, relative humidity, wind speed and direction, and solar irradiance are monitored. Using regression analysis, coefficients are obtained that represent building thermal performance characteristics of interest, HTC which constitutes a combined transmission and ventilation heat loss (Bauwens and Roels 2014).

The influence of transient effects induced by charging and discharging of the building's thermal mass can be reduced by sensibly choosing the experiment period and averaging the collected measurement data over a sufficient time span, achieving a quasi-steady-state approximation. The method requires long monitoring periods and uses aggregated data of at least 24 hour periods (or multiples thereof) as opposed to selected night time periods (Stamp 2015).

The co-heating test represents the state of the art with respect to measuring whole building fabric performance - it is the most comprehensive and accurate protocol available and is capable of achieving an uncertainty as low as $\pm$10% provided best-practice data collection and analysis methods are followed (Jack 2015). However, it requires typically 1-3 week monitoring during which the house must be unoccupied, and performing the monitoring is a labour, time, and skill intensive process (Stamp 2015) - and hence costly, limiting its application. Overcoming this obstacle is therefore key to enabling routine fabric performance measurement of buildings, enabling the sort of quality assurance measures and information provision that are likely to ultimately close the performance gap and reduce energy use for heating in domestic buildings (Jack 2015).

Most current co-heating trials have been performed on unoccupied houses, as a means of limiting uncertainty and confounding factors. Lowe and Gibbons (1988) remark the presence of occupants can mean variations in internal temperature and heat flows into and from the thermal mass can no longer be regarded as negligible. Earlier studies did monitor occupied houses however. As part of the Pennyland project, Chapman, Lowe, and Everett (1985) monitored 80 houses using 6 day aggregations, the HLC estimates showed reasonable correlation to the results from unoccupied dwellings. R. Everett, Horton, and Doggart (1985b), as part of the same study, concluded that this monitoring strategy could distinguish houses to within three basic levels of insulation (although not achieve an accuracy as good as $\pm$10%).

2.5.5 Modelling performance from energy metering

There are many diverse approaches to residential energy consumption modelling. Swan and Ugursal (2009) reviewed a range of sector-level analysis methods, choosing to divide these into top-down methods (econometric and technological) and bottom-up approaches, which they classify as being either engineering models or statistical models. However, the division between engineering and statistical models is not strict. In a review paper, Coakley, Raftery, and Keane (2014) identifies a spectrum of approaches to building energy modelling approaches from 'white-box' to 'black-box', via 'grey-box' Figure 2.5. Broadly, white-box models are derived from physical first principles, for example through a component-based bottom up approach. These tend to require extensive system information. Black box models on the other hand are purely empirical, making use of generic mathematical models or statistical tools to relate empirical inputs to outputs. Grey-box models sit between these extremes, taking advantage of domain knowledge to define a simplified model form whose parameters can be inferred from empirical data. In this case, knowledge of the thermal transfer processes in buildings is used to define an energy balance equation, and empirical data from smart meters and weather is used to infer the corresponding parameters in the energy balance equation.

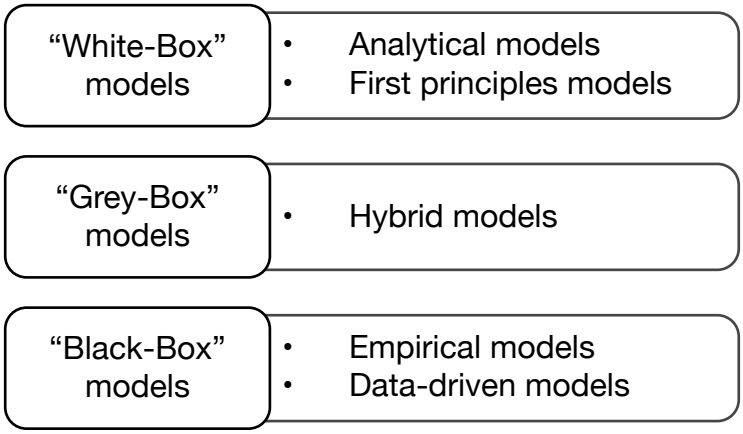

Figure 2.5: High level overview white, grey, and black box models. Moving from white to grey to black box models shifts the focus from analytical first principles models to empirical ones, with corresponding trade-offs in physical interpretability and analytic complexity.

In the building energy modelling domain, black box models benefit from short development times as they use existing statistical models or tools. Black box models - notably neural networks - have had success in pure demand prediction exercises (Edwards, New, and Parker 2012). However, there is no explicit link between black box model inputs and parameters and physical properties or observables. In addition to failing to provide a physical understanding of the system, it makes it impossible to calculate the expected effect of changes of those parameters, such as changing insulation levels. Furthermore, black box models tend to require extensive training data (Coakley,

Raftery, and Keane 2014).

Full building simulation tools such as EnergyPlus or TRNSYS may be considered as white-box models, as these use detailed physical descriptions of building components to calculate energy flows (Crawley et al. 2001). However, due to their level of detail they require very large amounts of input data. At best this can make them onerous to apply to existing dwellings, at worst this may expose them to the same problems of subjective inspections that affect EPC assessments.

Grey-box approaches can be broadly divided into steady-state methods and dynamic methods. Steady-state methods aim to simplify analysis by ignoring inter-period correlation and various transient terms, such as building thermal mass. It has been demonstrated that since the driving forces of dwelling energy consumption are quasi-periodic, choosing the correct time interval for analysis eliminates transient and higher-order terms from the model equations, reducing uncertainty and data requirements (Rabl 1988). However, the need to average data over long periods of time results in long data collection periods. Dynamic methods aim to capture all the time-dependent processes and can allow analysis using short monitoring periods.

An overview of steady-state methods and dynamic methods will now be given, followed by a discussion of their relative merits.

2.5.5.1 Dynamic methods

Dynamic methods use time-dependent models to take into account time dependent and time-lagged effects, such as the storage and release of thermal energy in a building's thermal mass. Full building simulation tools such as TRNSYS use dynamic models, however these white-box models rely on extensive detailed building information as inputs.

An extensive trial was conducted as part of IEA Annex 58 'Reliable building energy performance characterisation based on full scale dynamic measurements' (Roels 2012) in collaboration with the DYNASTEE project to use dynamic thermal modelling methods to analyse a heavily instrumented test cell under a variety of conditions. The test cell consisted of a transportable room with a single opening which supported a variety of configurations, such as various types of glazing, wall types, or insulation. The room was heavily instrumented with sensors internally and externally, and the experiment also collected extensive local weather data using a mobile weather station. The test cell was installed a several locations around Europe and the data generated used to inform and validate a range of dynamic models which formed the bases of the Annex final report. Measurements from this room was used to explore dynamic modelling methods. An extensive report on the Annex findings can be found in Janssens (2014).

Dynamic methods have been extensively applied to building components. Models applied include, autoregressive models RC-equivalents, transfer functions, and Bayesian statistical approaches.

These have been evaluated on both test cells and using in-situ measurements (Baker and Dijk 2008, Jiménez, Porcar, and Heras (2009), Biddulph et al. (2014), Bacher and Andersen (2014)). While these component-base methods are not always readily adaptable to whole dwelling performance analysis, they illustrate some of the key characteristics of dynamic methods. Advantages of dynamic methods as described by Baker and Dijk (2008) include:

- Generally shorted monitoring periods are required.
- Information about time-dynamic properties, such as thermal capacitance, may be inferred.
- Results may be more accurate.
- It may be easier to separate out different thermal processes, such as transmission from solar gain, or to model more processes (such as transmission to sky or ground) and in more detail.

Dynamic methods may also be more precise and accurate in cases where average heat flux density or temperature difference is small (Janssens 2014).

Rabl (1988) demonstrated in general the equivalence between time dependent thermal flow equations for whole buildings derived from first principles and Autoregressive–moving-average model with exogenous inputs (ARMAX) models, indicating that this family of models, which has been applied to building components, might also be successful for whole dwellings. The authors however found a relatively poor model fit, but suggested this may have been due to data collection issues. It has been noted however that ARMAX models as an 'off-the-shelf' statistical model do not identify physical information (Soderstrom et al. 1996). Instead, using the aforementioned equivalence it is possible to associate an ARMAX model parameter with a physical building property.

Juhl et al. (2013) developed the Continuous Time Stochastic Modelling for the R programming language (CTSM-R), a free open source tool for identifying physical models using real time series data. This was particularly suited for identifying the parameters of dwelling thermal-equivalence networks using monitored building data, enabling the estimation of lumped dwelling thermal parameters. Bacher and Madsen (2011) made use of this tool to identify which of a set of thermal-equivalence circuits best fit 5-minute interval data from an experimental housing lab called Flex-House - a single storey building test platform in Denmark equipped with internal temperature, local air temperature, solar irradiance, and heater input sensors. By using likelihood ratio tests, the best fit RC circuit for the test house was defined as in Figure 2.6. This demonstrated the possibility of identifying relatively complex thermal models for dwellings from monitored data.

Dynamic models have been used to develop load forecasting models. Bacher (2012) developed heat load forecasting models using adaptive linear time series models as a function of temperature, solar irradiance, and wind speed. This used a combination of data from previous time steps and Fourier series to model diurnal patterns, and was found suitable for time horizons up to 42 hours. Braun and Chaturvedi (2002) developed a grey box model of transient building heating and cooling load prediction, using RC-equivalence thermal network. Extensive building characteristic information

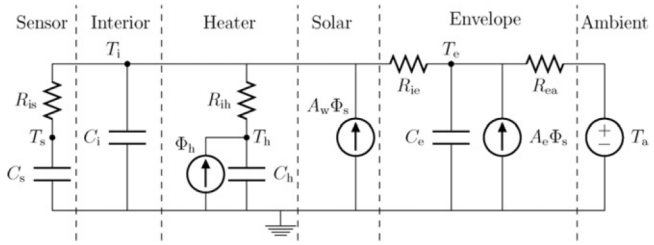

Figure 2.6: The final selected model TiTeThTsAe with individual model parts indicated (Bacher (2012))

and 2-4 weeks of operational data were required. The model was able to reproduce the building cooling load with RMS errors under 5%, but was not tested for forecasting.

Dynamic models are by definition more complex than their steady-state equivalents, as the latter are generally derived from the former by using simplifying conditions to eliminate terms in the model equations (Rabl 1988, Baker and Dijk (2008)). The data requirements for dynamic models are also more stringent - as illustrated by the attention required in the experimental protocol and the data validation recommended in the context of the PASLINK and DYNASTEE project (Baker 2004, Baker and Dijk (2008), Roels (2012), Janssens (2014)). Dynamic models may be more sensitive to noise and errors in input data relative to steady-state ones, since averaging data to obtain a steady-state approximation will tend to smooth the signal. When applying dynamic models, extensive pre-processing may be required such as applying suitable outlier and high-pass/low-pass filters (Bacher 2012). This process tends to require significant manual labour and attention and can be difficult to scale to large numbers of sites.

2.5.5.2 Steady-state methods

The co-heating approach discussed previously falls in the steady-state model category. This section will consider other examples of this type of analysis, with particular focus on the use of metered energy demand of occupied dwellings.

An early example of an application of steady-state methods to whole building analysis is the PRInceton Scorekeeping Method (PRISM), which used monthly billing data (commonly available in the USA) together with heating degree-day (HDD) data for the corresponding months to estimate annual energy demand (Fels 1986). A simple thermal model was defined:

$$\begin{aligned} P_{tot} &= P_B + \beta(T_h - T_{ex}) \text{ if } T_h - T_{ex} \geq 0 \\ P_{tot} &= P_B \text{ otherwise} \end{aligned} \tag{2.1}$$

where:

- P_{tot} (kW) is the total daily mean power demand, derived from monthly billing data.

- P_B (kW) is the base-load power demand.
- β (kW/$^\circ$C)is the house's effective heat loss rate.
- T_h ($^\circ$C) is the change point temperature between a base-load regime where mean power demand is a constant P_B and a heating regime where power demand is inversely proportional to outdoor temperature, related by β .

This model is ideally suited to the use of daily data, however at the time only monthly billing data was available. Consequently, monthly energy data was adapted with HDD data to fit the parameters of the model. Fels found that when using monthly data, the individual parameters of the fit had too much uncertainty and therefore defined a Normalised Annual Consumption (kWh) (NAC) (eq. 2.2) which estimated the total consumption for an average year using average HDD to the base T_h from 1970 to 1981 ($H_0(T_h)$). The NAC enabled reliable ranking of building performance, but due to the high uncertainty in the separate model parameters these could not be reported independently and could not be considered to be physically meaningful.

$$NAC = 365 * P_B + \beta * H_0(T_h) \tag{2.2}$$

PRISM has the advantage of being easy to understand and apply, an important feature for an energy performance index. A number of publications expanded on the original PRISM work. The American Society of Heating, Refrigeration, and Air Conditioning Engineers (ASHRAE) handbook includes extensions to PRISM to deal with different consumption patterns such as dwellings with air conditioning, introducing additional parameters to do so (ASHRAE 2002). However, the handbook advises against use of these models if possible because increasing the number of parameters to fit reduces the usefulness of the data (due to under-determination). Goldberg and Fels (1986) note that while the errors in the fit parameters β, P_B, and T_h are large compared to the 3-4% typical for NAC, when averaging over a large number of houses the individual parameter can offer valuable insight into changes in demand profiles. A method was developed to separate components of change in NAC over time and interpret the relationship between changes in these parameters and changes in building fabric. This was applied to monthly billing data from 243 gas heated dwellings in Wisconsin, USA. Although the spread of results was quite large, the median values were stable and informative. It was found that mean indoor temperatures, fabric losses, and base-loads did not change significantly over time. However, the large interquartile range of the change distributions of each parameter suggested that some individual dwellings did experience large changes in thermal properties, notably internal temperatures. This indicates that while aggregate analysis is beneficial to reducing parameter uncertainties, important information about individual dwellings is lost. Ruch and Claridge (1993) explored variants of PRISM using daily data from a range of commercial buildings. Rigorous error and goodness-of-fit statistics were developed to allow robust comparison between model variants - although these measures have not been taken up in subsequent literature. An interesting observation is that traditional fit measures such as

Root-Mean-Square Error (RMSE) do not take into account how well the fit performs over a range of temperatures.

The utility of PRISM in the UK context was until recently very limited, because energy readings in the UK are made irregularly in the domestic sector (generally once a year). In general, relatively little use appears to have been made to-date of PRISM outside of an original set of papers in the late 80s and early 90s. The introduction of smart meters in the UK has made such approaches relevant as suitable energy data is now being collected. Summerfield, Oreszczyn, et al. (2015a) used daily averages from smart meter data to calculate a Power Temperature Gradient (PTG) analogous to the heat loss rate β from PRISM. The PTG was used as a steady-state estimate of household energy performance without assigning physical meaning to it. Additionally, the external temperature above which heating is deactivated T_h is set to a building sample average value (about 15 $^\circ$C) rather than being estimated per site. This reduces the number of independent variables between buildings to facilitate comparison, however since the heating activation temperature has an important impact on energy consumption and is itself of considerable interest it would seem useful to also estimate this parameter during the model fit.

Bacher and Madsen (2013) modelled the district heating demand of 26 dwellings as a stationary function of ambient temperature, solar irradiance, and wind. Although the sample size was relatively small, this demonstrated the possibility of dwelling parameter estimation using smart meter and weather data only, without internal monitoring. A number of questions were left open by this research, including the suitably method to only consider days when heating was used, and the suitable model for wind direction.

Fels (1986), Rabl (1988), Bacher and Madsen (2013), and Summerfield, Oreszczyn, et al. (2015a), as well as the co-heating method, take similar approaches in terms of expressing thermal energy as linearised functions of weather parameters, with optionally additional internal dwelling information such as internal temperature. However, with regards to using dwelling energy meter data without a full on-site experimental protocol like the co-heating test, no canonical dwelling model or data analysis approach emerged from the review.

Steady-state analysis for building energy has also been applied in somewhat different contexts than whole-dwelling heat demand. Steady-state methods have been applied in the case of test cells, but have not been recommended for this case. The controlled nature of data collected in a test cell implies that input data is of sufficient quality for full dynamic modelling, therefore there are no particular advantages of limiting analysis to steady-state only (Soderstrom et al. 1996, Baker and Dijk (2008)). Baker and Dijk (2008) describes using steady-state methods to estimate the thermal transmission coefficient of building components in a test cell using a simple heat balance equation. The authors note that performing the type of analysis is very simple for a test cell where the element temperatures and thermal energy supplies, as well as external conditions, are well known. However, longer averaging periods (and therefore monitoring periods) are required to reach steady-

state conditions, and the dynamic properties (such as thermal masses) of the elements cannot be determined. Birt et al. (2012) disaggregated electrical consumption of Canadian dwellings using piecewise linear functions. This work included a thermal component due to the high cooling demand in summer of the buildings studied. However no dwelling physical properties were inferred.

2.5.5.3 Comparison of dynamic and steady state methods

Table 2.1 summarises the method type, building type, number of buildings, and monitoring period for the papers discussed in the dynamic methods and steady state-methods sections. Only papers which explicitly discussed application to field trial data are considered. The key characteristics of dynamic and steady state models in the literature relevant to modelling dwelling thermal performance from smart meter data are presented in Table 2.2.

It is immediately apparent from Table 2.1 that research based on steady-state methods has been applied to larger numbers of dwellings, but has generally required much longer monitoring periods. More model inputs are used for dynamic methods, notable internal temperatures and in several cases measurements of the heating energy supply (rather than the total energy supply). Additionally, as noted in Table 2.2, the additional complexity and requirements on data quality may make the application of dynamic methods prohibitively costly in terms of manual labour as the number sites studied grows.

Table 2.1: Comparison of the addressed building types, number of buildings studies, and monitoring period for the papers reviewed in the dynamic methods and steady state-methods sections, considering only papers discussing experimental trials.

Reference	Method type	Building type	N buildings	Monitoring period	Model inputs
Baker and Dijk (2008)	Dynamic	Test cell	1	10 days	Internal & external temperature, heat flux measurements, solar irradiance, wind speed, heating energy demand
Jiménez, Porcar, and Heras (2009)	Dynamic	Test cell	1	24 days	Internal & external temperature, heat flux measurements, solar irradiance, wind speed, heating energy demand

Reference	Method type	Building type	N buildings	Monitoring period	Model inputs
Bacher and Andersen (2014)	Dynamic	Test cell	1	10 days	Internal & external temperature, heat flux measurements, solar irradiance, wind speed, heating energy demand
Janssens (2014)	Dynamic	Test cell	1	1+ weeks	Internal & external temperature, heat flux measurements, solar irradiance, wind speed, heating energy demand
Biddulph et al. (2014)	Dynamic	Solid wall house	1	14 days	Internal & external temperature, heat flux measurements, solar irradiance, wind speed, heating energy demand
Rabl (1988)	Dynamic	Office building	1	14-19 days	Internal & external temperature, heat flux measurements, solar irradiance, wind speed, heating energy demand, air exchange rate
Bacher and Madsen (2011)	Dynamic	Test house	1	6 days	Internal & external temperature, solar irradiance, wind speed, heating energy demand
Bacher (2012)	Dynamic	Houses	16	2.5 years	External temperature, solar irradiance, wind speed, heating energy demand

Reference	Method type	Building type	N buildings	Monitoring period	Model inputs
Fels (1986)	Steady state	Houses	1000	1 year	External temperature, heating degree days, monthly energy bill
Goldberg and Fels (1986)	Steady state	Houses	243	1 year	External temperature, heating degree days, monthly energy bill
Bacher and Madsen (2013)	Steady state	Houses (detached)	26	2 years	External temperature, solar irradiance, wind speed, district heating energy demand
Summerfield, Oreszczyn, et al. (2015a)	Steady state	3-bed houses	255	2 years	External temperature, gas energy
Chapman, Lowe, and Everett (1985)	Steady state	Semi-detached dwellings	80	2 years	Internal & external temperature, electricity & gas energy, solar irradiance

Table 2.2: Overview of key differences of dynamic and steady state building thermal models as relevant to grey-box modelling of dwelling thermal performance

Aspect	Dynamic	Steady state
Model complexity	Complex	Simple
Data resolution	Hourly or higher	Daily or lower
Data quality required	High (Clean continuous time series)	Low (Robust to errors and missing data)
Data collection duration	Short (days/weeks)	Long (weeks/months)
Manual labour required per site	Extensive	Minimal
Dynamic properties inferred	Yes	No

Overall, steady state methods appear to better fit the particularities of smart meter data. Smart meter infrastructure facilitates the collection of energy readings over long periods of time, but offers little in terms of quality control or the ability to check a given dwelling to help address issues that might appear during data collection. As a result it is important for the analysis method to be robust to data quality issues. Additionally, the goal of applying the model in an automated manner to large numbers of sites reinforces the need for a robust approach, since it would be prohibitive in terms of manual labour to have to individually inspect and correct sites where the analysis failed. Steady-state methods emerge as being significantly simpler than dynamic ones. This advantage becomes more pronounced as the number of dwelling to analyse increases and the quality of data collection decreases.

By contrast, dynamic methods appear to be better suited to building monitoring campaigns, as the reduced number of dwellings combined with the possibility to check monitoring equipment on-site addresses the data quality issues and implies that increase manual labour for data collection and processing is not a significant drawback. Additionally, the ability to estimate transient properties (such as thermal mass) and do so using much shorter monitoring periods are beneficial when there are limited resources available to run the monitoring (install, check, and later remove sensing equipment).

2.6 Summary

This chapter has provided an overview of the UK residential energy demand and energy efficiency policies across EU and UK. The importance of assessing dwelling energy performance has been highlighted.

It has been argued that RdSAP assessment method, and by extension the EPC, suffers from a number of issues that undermine it's suitability for delivering dwelling energy performance information. Crucially, its dependence on manual dwelling inspections implies that significant scope exists for human error in the assessments which undermine potential improvements to the underlying model. This motivates the search for a dwelling assessment approach that uses measured dwelling data.

In-situ U-value measurements, infra-red thermography, and co-heating experimental measuring approaches were considered. U-values were found to be measurable to high degrees of accuracy, but require careful experimental implementation with access to the dwelling, and can only measure individual elements rather than the whole building envelope. Infra-red thermography enables rapid qualitative assessment of dwelling envelopes but is limited in its ability to provide quantitative information. Co-heating was found to provide the most comprehensive assessment of dwelling envelopes, but at the cost of an onerous experimental protocol requiring empty dwellings with intensive monitoring.

An overview of smart metering in the UK was presented. These meters are intended to bring a number of benefits, among which is that of being a reliable source of energy data for the purpose of developing new approaches to assessing and understanding energy demand. They should enable data for large numbers of dwellings over long periods of time to be collected.

Grey-box approaches to building modelling were introduced. These approaches combined aspects of bottom-up physical models with established statistical tools to infer building properties for measured data, and as such appear a promising approach for developing dwelling energy models using smart meter data.

A selection of studies discussing dynamic and steady state models for the inference of thermal properties of buildings from monitored data were presented. The advantages and disadvantages of each approach were contrasted. Dynamic models were found to require short monitoring periods (on the order of a few days to a few weeks) and were generally performed with larger numbers of sensors (both internal and external) recording at short intervals (of the order of 5-10 minutes). They enabled inference of dynamic characteristics such as building thermal mass. However, dynamic models were found to be considerably more complex that steady state models. In particular, they were found to require high quality input data, which implied extensive manual labour to obtain, clean, and error check. By contrast, steady state methods were relatively robust to such issues. This is reflected in the sample size differences between the dynamic and steady state publications, where the dynamic methods were mostly applied to single buildings while steady state methods could be applied to hundreds or thousands of dwellings. Steady state methods therefore should scale more readily to large numbers of dwellings, which makes them particularly interesting for this research.

The methods reviewed shared common theoretical underpinnings based on the dwelling heat balance. The heat balance could be described in more or less aggregated ways depended on the level of detail required, some studies developed detailed RC-equivalence thermal networks while others (notably the steady state ones) lumped all parameters into a small number of building specific coefficients. In general dynamic methods were able to infer more dwellings properties, both constant and dynamic, but to do so made use of more detailed buildings descriptions and model inputs. Steady state models were sufficient for estimating lumped parameters.

Steady state models were found to be well suited to the data collection profile of smart meters, in that these meters facilitate bulk collection of energy data over long periods for large numbers of dwellings but have limited scope for ensuring data quality or controlling the conditions (e.g. weather conditions) under which the data is collected. Steady state models are more robust to data issues, and have been applied to large datasets in the past. However, although some research was published using smart meter data for thermal dwelling performance analysis, no established methodology was identified (in contrast to the co-heating test for example).

It seems therefore reasonable to conclude that a grey-box steady state approach should be suitable to infer the thermal properties of dwellings from their smart meter data, and to be able to do so 'at scale' (i.e. on large numbers of dwellings). No established model and model fitting approach exists to perform this analysis. This research will therefore in the following chapters present a dwelling thermal model suitable for use with smart meter data, and a model fitting and data processing approach to enable robust parameter estimation for large numbers of dwellings.

Chapter 3

Theory

3.1 Introduction

This chapter defines building physics terms and concepts including key heat transfer processes, namely the conduction, radiation, convection, air exchange and internal gains, that are relevant to calculating a building's thermal performance. Equations for dwelling net energy balance and the parameters required to describe this are presented based on whole-dwelling approximations from literature. An equation for the whole dwelling heat balance is given, and related to the metered energy consumption (Section 3.9). Based on the metered energy-heat balance relation, expressions are derived for the whole dwelling Heat Transfer Coefficient (HTC). A simple model for dwelling internal temperature is given (Section 3.4). Using this model and the expression for HTC, a whole dwelling Heat Loss Coefficient (HLC) is derived, thereby operationalising the thermal performance through the metered energy demand.

3.2 Thermal energy balance

The dwelling net heat flow rate Φ_{tot} (kW) is defined using established formulations of dwelling heat transfer processes. Energy flowing out of the building is defined as negative, while thermal energy flowing into the build is positive. The thermal model calculates a thermal balance as the sum of contributions to heat flow in/out of the dwelling Φ_i:

$$\Phi_{tot} = \sum_{i}^{n} \Phi_i \tag{3.1}$$

Linking functions are then defined which describe how metered energy demand P_{tot} (kW) changes as a function of the thermal balance eq. 3.2. This enables the operationalisation of dwelling thermal

performance through measured energy demand, while making assumptions explicit through the mathematical formulation. The heat loss coefficient (kW/K), which represents the combined heat loss rate of the dwelling, will be defined in terms of the component thermal flows.

$$P_{tot} = f(\Phi_{tot}) \tag{3.2}$$

A model of dwelling internal temperature is defined, as there are many cases where internal temperature is not measured.

This research aims to maintain a clear distinction between energy and power. Calculations are performed in units of mean power demand (kW) over suitable intervals. Yearly energy consumption (kWh/year) is also used where appropriate, principally for comparison with other demand estimates. Computer code and calculations are performed in Kelvin, however values are reported in this thesis in $^\circ$C for consistency with existing literature and to assist the intuitive understanding of reported results. Note that per-degree values (such as HTC) are entirely equivalent in either unit of temperature (kW/K and kW/$^\circ$C). The $^\circ$C value will be given for consistency.

3.3 Steady-state approximation

Determining thermal flows normally requires solving dynamic heat flow equations and taking the thermal capacity of building elements into account. The sum of the heat capacity of building elements is approximated as the building's thermal mass (eq. 3.3):

$$Q = C_{th}\Delta T_m \tag{3.3}$$

where Q (kJ) is the thermal energy transferred, C_{th} is the effective thermal mass of the combined dwelling elements, and ΔT_m is the change in effective temperature of the combined thermal mass elements. The thermal mass may act as a heat source or a heat sink depending on the dwelling internal temperature and the energy stored in the thermal mass (measured through its effective temperature). Calculating the net energy balance at a given time requires considering the energy stored in preceding time steps.

However, the drivers of heat loss and gain in real buildings are highly periodic across days, driven by the diurnal cycle of weather and accompanied by set heating schedules. As a result it is possible to obtain a steady-state approximation of the energy flow that does not depend on the thermal mass by choosing an appropriate periodic average, as these terms cancel out when suitable boundaries are chosen for the integration (Rabl 1988).

3.4 Internal temperature model

Dwelling internal temperatures are not widely monitored compared to the number of dwellings with smart meters. Energy losses through the dwelling fabric and through ventilation are driven by the internal temperature relative to the external temperature. It is therefore necessary to model internal temperatures where they cannot be measured. The simplest assumption is for mean steady-state internal temperatures to be constant, as in (Fels, Goldberg, and Lavine 1986, Summerfield, Oreszczyn, and Hong (2013)), or similarly that internal temperatures vary about a mean value but are not covariant with the external temperatures. However, Summerfield et al. (2007) demonstrated on a sample of 14 gas-heated dwellings that there was significant covariance between internal and external temperatures.

Furthermore, it can be demonstrated analytically using a simple thermal model (see Appendix A) that, given a typical heating pattern, the mean internal temperature varies with external temperature, as the internal temperature will decay towards the external temperature plus a temperature correction due to internal gains during the non-heated periods. In other words, internal temperature is affected by external temperature due to the change in base internal temperature. The SAP model accounts for this relation, however it does so in a way that is not convenient for data-driven modelling (BRE 2011). A number of models aim to predict mean internal temperatures, generally as a function of building types and socio-economic factors (Huebner et al. 2014), however these models do not capture the relation between daily external and internal temperatures.

As a first order approximation, a linear model for internal temperature is defined as a function of external temperature during the heated days. Love (2014) presents evidence from BREDEM modelling suggesting that a linear model is appropriate when heating schedules are constant. Furthermore, empirical work in Section 6.2 demonstrates that a linear approximation is reasonable for a majority of sites in a representative sample of English dwellings. The model is given by eq. 3.4

$$T_{in} = C_T(T_{ex} - T_{fix}) + T_0 \tag{3.4}$$

where C_T is the dimensionless covariance of T_{in} ($^\circ$C) with T_{ex} ($^\circ$C) (i.e. $\frac{dT_{in}}{T_{ex}}$) and T_0 ($^\circ$C) is the mean internal temperature at external temperature $T_{fix} = 5\,^\circ$C. Setting T_{fix} enables standardisation of the temperature model and enables cross comparison between dwellings and regions. A value of 5 $^\circ$C is chosen in order to facilitate comparison with existing work in this field including Hamilton et al. (2016b), Oreszczyn et al. (2006), and Wilkinson et al. (2001). Hamilton et. al investigated the effect of standardising to higher (10 $^\circ$C) and lower (0 $^\circ$C) values and found no notable effect on study outcomes, the choice of 5 $^\circ$C is therefore reasonable albeit somewhat arbitrary. The choice of standardising to 5 $^\circ$C produces more intuitive results compared to using the y-intercept in degrees-Kelvin, and could be considered slightly more intuitive than 0 $^\circ$C or 10 $^\circ$C.

3.5 Fabric heat transfer

Conduction is the main heat transfer process through building fabric elements. The rate of energy transfer $\Phi_{tr,i}$ (kW) by conduction across a solid element i is a function of the area A_i, and the element U-value, which encapsulates the element thickness, conductivity, and skin properties regulating the transfer of energy from the element surface to the surrounding air (eq. 3.5). The U-value is a commonly used empirical measure of a build element's insulating properties and is usually approximated as a constant, although it is recognised that in reality it may vary according to some environmental conditions (humidity, wind speed).

$$\Phi_{tr,i} = UA_i(T_{ex} - T_{in}) \tag{3.5}$$

The total dwelling fabric heat transfer is the sum over the individual elements, which can be approximated by the whole dwelling heat transfer coefficient H_{tr} (ISO 2010). This coefficient combines direct heat transfer from internal conditioned space across the fabric to the external environment, as well as heat flows to the ground, to unconditioned spaces, and to adjoining buildings. The total conductive heat transfer Φ_{tr} is given by eq. 3.6 as a function of this coefficient and the internal-external temperature difference.

$$\Phi_{tr} = H_{tr}(T_{ex} - T_{in}) \tag{3.6}$$

Note that this formulation makes the simplifying assumption that there is a single value for internal and external temperatures. This means assuming a single mean internal temperature for the building rather than attempting to model multiple thermal zones. It also assumes the external environment has one temperature - this would not be the case for heat transfer to the ground, which is likely to have a different temperature than the ambient air (Stamp 2015).

3.6 Radiative heat transfer

3.6.1 Solar gains

Solar irradiance is an important contributor to dwelling thermal balance (Orehounig, Dervishi, and Mahdavi 2011). Energy is absorbed by a building both through the heating of opaque surfaces and via transparent windows that transmit and retain radiation inside the building. The total energy absorbed is a function of the number, area, orientation, and optical and thermal properties (absorptivity, radiative heat transmission) of the individual building elements.

Bauwens and Roels (2014) describe the linearisation in eq. 3.7 used in co-heating tests:

$$\Phi_{sol} = A_{sol} I_{sol} \tag{3.7}$$

where A_{sol} in m^2 is the effective solar aperture of the building and I_{sol} is the received solar irradiance (kW/m^2). This is a convenient simplification which was adopted for this research. The relation between the whole dwelling aperture and the solar gain through individual elements is further described in Appendix B. It is important to note that the effective solar aperture cannot be simply linked to a single measurable building characteristic, but is rather an index of the propensity of a dwelling to gain energy from solar irradiance.

3.6.1.1 Solar irradiance on arbitrary surfaces

The irradiance I_{sol} input data should be chosen to best represent the mean irradiance received by building surfaces. In co-heating experiments irradiance on a vertical south-facing plane was used (R. Everett, Horton, and Doggart 1985a; Bauwens and Roels 2014; Stamp 2015; Jack 2015). Meteorological data sources quote irradiance figures as Global Horizontal Irradiance (GHI) (kW/m^2), being the total irradiance (direct plus diffuse irradiance) per unit area on a horizontal plane. A number of models enable the estimation of irradiance on different orientations using GHI data and other data on the state of the atmosphere, notably making it possible to split GHI into direct and diffuse components, and then re-project those onto an arbitrary surface as a function of the location of the surface on the earth and the position (azimuth and zenith angles) of the sun. The input parameters are:

- The solar position in the sky at the site location, using a Solar Position Algorithm (SPA) National Renewable Energy Laboratory (NREL).
- The extraterrestrial Direct Normal Irradiance (DNI-E) (which can be derived from solar position, and site location).
- The surface level Direct Normal Irradiance (DNI) as a function of GHI, dew-point temperature and air pressure.
- Diffuse Horizontal Irradiance (DHI) as a function of apparent solar zenith angle.
- Total irradiance I_s as a function of DNI, GHI, DHI, atmospheric conditions, and ground properties.

Table 3.1 lists the method selected for the calculation of each of the above parameters. In order to improve the accuracy of the DIRINT method listed, dew-point temperature and air pressure data may be used in addition to the GHI to improve the algorithm's estimate of the diffusion of solar radiation through the atmosphere. The total irradiance model is based on the work of Loutzenhiser et al. (2007), who compared 7 established models (Klucher, Hay–Davies, Reindl, Muneer, Perez 1987, and Perez 1990) and found the Perez 1990 model to be the best performing by a small margin. The PVLIB Python module developed by the Sandia National Laboratory was

used to provide peer-reviewed implementations of the parameter estimation methods (Andrews et al. 2014).

Table 3.1: Selected method and reference for estimation of solar parameters

Parameter	Method
Solar Position	NREL SPA (Reda and Andreas 2004)
DNI-E	Spencer, in (Reno, Hansen, and Stein 2012)
DNI	DIRINT (Ineichen 2008)
DHI	GHI - DNI * cos(θz)
I_s	Perez 1990 (Loutzenhiser et al. 2007)

3.6.2 Longwave radiative losses

Longwave (infra-red) radiation is generally considered in terms of losses because at temperatures typical of building envelopes bodies emit mainly in the longwave spectrum. For the purposes of this work, it was assumed that longwave radiative losses would not be sufficiently important to consider, on the basis of work by Stamp (2015). Exploratory calculations found that for a range of sky conditions the heat loss should in any case be small. Furthermore, estimation of sky temperatures and building skin emissive temperatures were considered to be likely to introduce excessive uncertainty. An equation for the radiative loss to the sky is given in Appendix B.

3.7 Air exchange

Ventilation heat loss, through the exchange of warm internal air for cold external air, is thought to account for approximately 20% of a dwelling's heat loss (Cooper and Palmer 2013). Outside airflow into the building is mostly driven by pressure differences caused either by the winds or resulting from the difference between inside and outside temperatures. The heat loss due to infiltration is given in terms of volume flow rate q_ve (m^3/s) and density ρ (kg/m^3) (1.205 kg/m^3 at standard temperature and pressure) (eq. 3.8) and is due to the energy content of the exchanged air, defined by its heat capacity c_p (J/kgK) (1005 J/kgK at standard temperature and pressure).

$$\Phi_{ve} = q_{ve} c_p \rho (T_{ex} - T_{in}) \tag{3.8}$$

The air flow q_ve is driven by the pressure difference between inside and outside the building induced by the stack effect (flow of air through a building resulting from a pressure differential created

by the difference in temperature between internal and external air) and wind force. An extensive theoretical treatment of these is given by Walker and Wilson (1998) and summarised in Appendix B. The resulting total ventilation heat exchange is given by eq. 3.9, with a base ventilation rate coefficient C_{ve}, a wind coefficient C_{wind} and a stack coefficient C_{stack}.

$$\Phi_{ve} = c_p \rho_{air} (T_{ex} - T_{in})(C_{ve} + C_{stack}(T_{ex} - T_{in}) + C_{wind} v_{wind}) \tag{3.9}$$

This is further linearised, analogously to co-heating tests as in Stamp (2015), which further simplifies the definition of the whole dwelling heat transfer coefficient (see Section 3.10). The linearised ventilation heat loss rate is given by eq. 3.10, which makes the simplifying assumption that the non-linear ventilation heat loss can be approximated by a linear term as a function of the internal-external temperature difference and the coefficient C_{ve}. This simplification is evaluated in Section 6.1.

$$\Phi_{ve} = c_p \rho_{air} C_{ve} (T_{ex} - T_{in}) \tag{3.10}$$

3.8 Internal gains

These gains are in the general category of unintentional gains - heat dissipated from sources other than the heating system and solar gains. Determining these gains requires first determining how much energy is dissipated by each source, then deciding how much of that energy will contribute to the building's heat balance. ISO 13790 considers internal gains from

- heat flow from occupants (metabolic gains)
- heat from appliances
- heat from lighting
- heat from hot water mains and sewage

Metabolic gains were not considered to be a significant contributor to dwelling energy balance, as these are of the order of 60W compared to heating system power of the order of several kW. Heat from appliances and lighting are encapsulated in the base-load power related gains, discussed below. Heat gains from hot water are also included in the treatment of base-load, as the energy imputed to the hot water supply is also metered through the total consumption.

3.8.1 Base-load and incidental gains

The gains from energy used in lighting, appliances, plug loads, as well as water heating is captured in the base-load parameter P_B. The base-load is the mean power outside the heating periods. The

incidental gains may have complicated determinants, for which we require a suitable approximation that can be related to measured demand. For this we introduce a parameter η_B describing the fraction of base-load power that contributes to the net internal thermal gains. The thermal contribution of base-load is:

$$\Phi_B = \eta_B P_B \tag{3.11}$$

where $0 \leq \eta_B \leq 1$. In this research we assume that $\eta_B \cong 1$, since almost all energy used in the dwelling will eventually dissipate as heat.

3.8.2 Lighting and Appliances

The heat gains from lighting and appliance use should be included in the calculation for base-load thermal gains. From a thermodynamic perspective, appliances that consume energy will eventually convert that energy to heat. In simple cases such as incandescent light bulbs energy is converted directly to heat. In more complicated cases such as washing machines, energy is converted into a combination of hot water and mechanical energy, which eventually dissipates as heat.

3.8.3 Water heating

Domestic hot water supply is important to consider because some metered energy will be lost down the drain without contributing to the dwelling thermal balance - although some will be recovered from water storage and piping (ISO 2010). The amount of hot water use is mainly determined by the number of occupants (Energy Saving Trust 2008).

Although important, it is likely that hot water may have less of an impact than traditionally expected. The Energy Saving Trust performed measurements of hot water demand across 120 houses equipped with a mixture of regular and combi-boilers (Energy Saving Trust 2008). Mean delivery temperature of domestic hot water was 52 $^\circ$C, significantly below the assumed value of 60 $^\circ$C, while inlet temperatures were on average 5.3 $^\circ$C higher than the assumed value of 10 $^\circ$C - so the boilers provided a temperature rise of 36.7 $^\circ$C instead of the assumed 50 $^\circ$C, resulting in significantly lower energy use for hot water. Daily mean power used for heating hot water was 0.124$\pm$0.0025kW (consumption of 120$\pm$2 litres/day). Given a national mean consumption of 2.2kW (DECC 2013a), this would represent at most a 5% error in total consumption if none of the heat was recovered within the dwelling.

3.8.4 Secondary heating systems

48% of UK dwellings use some form of secondary heating system (Hulme, Beaumont, and Summers 2013c). These are predominantly gas and electric heaters. Since the energy use of these heating systems is metered there is no need for special consideration for the dwelling energy balance. This however highlights the importance of using total metered energy demand to capture all heat sources. In the case of the approximately 10% of dwellings with solid fuel or 'other' supplementary heating systems however it is simply not possible to model their contribution to the thermal balance.

3.9 Metered energy demand

3.9.1 Relating thermal flow to boiler energy consumption

By definition in a steady-state model the energy balance equation must hold:

$$Q_{in} - Q_{out} = 0$$

Assuming that the heating energy is needed when the net heat flow is negative ($\Phi_{tot} < 0$), the heating power demand $P_{H,D}$ (eq. 3.12) is equal to the absolute value of net heat flow such that the steady-state heat balance is zero.

$$P_{H,D} = \begin{cases} |\Phi_{tot}| \text{ if } \Phi_{tot} < 0 \\ 0 \text{ otherwise} \end{cases} \tag{3.12}$$

The boiler and heating system efficiency impacts the relation between metered fuel consumption and delivered energy. In the case of electric resistance heating, there is generally no need to consider boiler efficiency as conversion from electricity to heat is lossless. This would not be the case for electric heat-pump systems which will have a coefficient of performance >1. For this reason heat-pumps are particularly challenging with regards to estimating thermal flows from metered demand, treatment of this topic is outside of the scope of this research.

Gas boilers lose energy through imperfect combustion, in the heat exchanger, and through flue gases. Condensing boilers aim to minimise the latter by extracting latent heat from the flue gas, reducing the flue outlet temperature. Orr, Lelyveld, and Burton (2009) conducted in-situ monitoring of condensing boiler efficiencies and found a mean efficiency of 82.5%, calculated as (space heating energy + hot water energy)/gas (similar to the boiler rating standard BS EN 483 and related BRE laboratory manual (En 2006; Shiret and (BRE) 2010)).

The power P_H required by the heating system to meet the heating demand power $P_{H,D}$ as defined in eq. 3.12 can be modelled using the heating system efficiency η_{HS}:

$$P_H = \frac{P_{H,D}}{\eta_{HS}} \tag{3.13}$$

for a single heating system. For multiple heating systems the efficiency adjusted supplied power must sum to the total heating demand:

$$\sum_i \eta_i P_{H,i} = P_{H,D} \tag{3.14}$$

In most cases it is only practical to consider a single η_{HS} which will be the mean efficiency in the case of multiple heating systems or the single boiler efficiency in the case of a single heating system. In theory, η_{HS} may also vary with external temperature - particularly in the case of heat pumps (Elwell et al. 2015). This variation is assumed not to be significant for the purpose of this thesis.

3.9.2 Conversion of gas volume to energy

Gas consumption may be recorded as a volume in m^3. Gas energy content in the UK is given by the National Grid official formula (NationalGrid 2015):

$$energy(kWh) = consumption(m^3) * CV(MJ/m3) * 1.02264/3.6 \tag{3.15}$$

where CV is the gas Calorific Value (MJ/m^3) defined at a standard temperature of 15 $^\circ$C (288.15K) and pressure of 101.325 kPa. The factor of 3.6 converts from MJ to kWh. Gas CV may vary over time and between gas distribution regions. The correction factor of 1.02264 accounts for the difference between the standard temperature and pressure used for quoting CV and the typical conditions at the domestic gas metering point.

3.10 Heat Transfer Coefficient

The Heat Transfer Coefficient (HTC) in kilowatts per degree Centigrade or Kelvin (kW/ $^\circ$C or kW/K) of a dwelling is used as an indicator of overall steady-state dwelling fabric thermal performance. It has been defined in co-heating tests as the absolute change in quasi-steady-state dwelling heat flux with change in temperature difference between the internal and external environments (eq. 3.16) (ISO 2010; Stamp 2015; Jack 2015; Bauwens and Roels 2014):

$$HTC = |\frac{d\Phi_{tot}}{d\Delta T}| \tag{3.16}$$

Where $\Delta T = T_{ex} - T_{in}$. The dwelling net heat flux Φ_{tot} is given by eq. 3.17:

$$\Phi_{tot} = \Phi_{tr} + \Phi_{ve} + \Phi_{sol} + \Phi_B \tag{3.17}$$

where

- Φ_{tr} : fabric transmission losses (negative when heat flows out of dwelling) as defined in eq. 3.6.
- Φ_{ve} : ventilation losses (negative when heat flows out of dwelling) as defined in eq. 3.9.
- Φ_{sol}: solar gains (positive heat gain) as defined in eq. 3.7.
- Φ_B : base-load/incidental gains (positive heat gain) as defined in eq. 3.11.

From eq. 3.7, $\frac{d\Phi_{sol}}{d\Delta T} = 0$, while the controlled conditions of the co-heating experiment imply that base-load or incidental gains are eliminated ($\Phi_B = 0$). From eq. 3.6 $|\frac{d\Phi_{tr}}{d\Delta T}| = H_{tr}$, therefore:

$$HTC = |\frac{d\Phi_{tot}}{d\Delta T}| = H_{tr} + \frac{d\Phi_{ve}}{d\Delta T} \tag{3.18}$$

The ventilation loss term is linearised such that HTC is a building constant and the sum of the fabric and ventilation losses is given by eq. 3.19.

$$\Phi_{tr} + \Phi_{ve} = HTC\Delta T \tag{3.19}$$

3.10.1 HTC and occupied dwelling demand

Combining the heating power demand equation eq. 3.12 with the boiler efficiency, eq. 3.13 and accounting for base-load power P_B gives the piecewise model for dwelling power demand eq. 3.20.

$$P_{tot} = \begin{cases} \frac{1}{\eta_{HS}}|\Phi_{tot}| + P_B \text{ if } \Phi_{tot} < 0 \\ P_B \text{ otherwise} \end{cases} \tag{3.20}$$

For the heating period therefore $P_{tot} = \frac{1}{\eta_{HS}}|\Phi_{tot}| + P_B$. Expanding the thermal balance Φ_{tot} using eq. 3.17 and taking the derivative with respect to the change in internal-external temperature difference gives eq. 3.21.

$$\begin{aligned}\frac{dP_{tot}}{d\Delta T} &= \frac{1}{\eta_{HS}}\frac{d\Phi_{tot}}{d\Delta T} + \frac{dP_B}{d\Delta T} \\ &= \frac{1}{\eta_{HS}}(\frac{d\Phi_{tr}}{d\Delta T} + \frac{d\Phi_{sol}}{d\Delta T} + \frac{d\Phi_{ve}}{d\Delta T} + \frac{d\Phi_B}{d\Delta T}) + \frac{dP_B}{d\Delta T}\end{aligned} \tag{3.21}$$

The base-load gains Φ_B are a function of P_B. Analysis performed in Section 6.1.1 suggests that it is reasonable to assume that base-load power demand is also not temperature dependent, i.e. $\frac{dP_B}{d\Delta T} = 0$, implying that $\frac{d\Phi_B}{d\Delta T} = 0$. Therefore differentiating P_{tot} with respect to change in ΔT gives eq. 3.22 as an equivalent of eq. 3.18 for an occupied dwelling.

$$\frac{dP_{tot}}{d\Delta T} = |\frac{1}{\eta_{HS}}\frac{d\Phi_{tot}}{d\Delta T}| = \frac{1}{\eta_{HS}}(H_{tr} + \frac{d\Phi_{ve}}{d\Delta T}) = \frac{1}{\eta_{HS}}HTC \tag{3.22}$$

The total power demand during the heating regime where the net thermal flow is negative is summarised in eq. 3.23.

$$P_{tot} = \frac{1}{\eta_{HS}}|\Phi_{tot}| + P_B \tag{3.23}$$

As $\Phi_{tot} < 0$, $|\Phi_{tot}| = -\Phi_{tot}$. From eq. 3.25 :

$$\begin{aligned}|\Phi_{tot}| &= |HTC(T_{ex} - T_{in}) + A_{sol}I_{sol} + \eta_B P_B| \\ &= -(HTC(T_{ex} - T_{in}) + A_{sol}I_{sol} + \eta_B P_B) \\ &= HTC(T_{in} - T_{ex}) - A_{sol}I_{sol} - \eta_B P_B\end{aligned} \tag{3.24}$$

The total power is therefore given by eq. 3.25.

$$P_{tot} = \frac{1}{\eta_{HS}}(HTC(T_{in} - T_{ex}) - A_{sol}I_{sol} - \eta_B P_B) + P_B \tag{3.25}$$

3.11 Heat Loss Coefficient

In eq. 3.18 it was demonstrated that the change in total power demand P_{tot} with change in internal-external temperature difference ΔT is related to the Heat Transfer Coefficient HTC and the heating system efficiency η_{HS}. However, inspection of eq. 3.25 indicates that it is not possible to separate HTC and η_{HS} without additional information as these parameters are covariant with respect to ΔT. A Heat Loss Coefficient (HLC) is therefore defined which incorporates thermal losses from the fabric and the heating system (eq. 3.26)

$$HLC = \frac{1}{\eta_{HS}}HTC \tag{3.26}$$

Substituting eq. 3.26 into eq. 3.25 gives eq. 3.27 describing the total power demand during the heating regime.

$$P_{tot} = HLC(T_{in} - T_{ex}) - \frac{1}{\eta_{HS}} A_{sol} I_{sol} - \frac{\eta_B}{\eta_{HS}} P_B + P_B \tag{3.27}$$

Outside of the heating regime the power demand is simply the base-load power P_B.

3.11.1 Power model using internal temperature model

The model described in eq. 3.27 depends on the difference ΔT between internal temperature T_{in} and external temperature T_{ex}. While T_{ex} may be derived from weather data using a site's location, T_{in} would normally require data recording from temperature sensors, which most dwellings are not equipped with. Substituting the linear temperature model eq. 3.4 ($T_{in} = C_T(T_{ex} - T_{fix}) + T_0$) into eq. 3.27 gives eq. 3.28.

$$\begin{aligned} P_{tot} = P_B + \tfrac{1}{\eta_{HS}}(HTC(C_T(T_{ex} - T_{fix}) + T_0 - T_{ex}) \\ - A_{sol} I_{sol} \\ - \eta_B P_B) \end{aligned} \tag{3.28}$$

Re-arranging eq. 3.28 to separate coefficients for T_{ex}, I_{sol}, P_B and T_0:

$$\begin{aligned} P_{tot} = -\tfrac{HTC}{\eta_{HS}}(1 - C_T)T_{ex} + \tfrac{HTC}{\eta_{HS}} T_0 \\ - \tfrac{HTC}{\eta_{HS}} C_T T_{fix} - \tfrac{A_{sol}}{\eta_{HS}} I_{sol} + P_B(1 - \tfrac{\eta_B}{\eta_{HS}}) \end{aligned} \tag{3.29}$$

Using the definition of HLC (eq. 3.26) to obtain a linear expression for the total power (eq. 3.30).

$$\begin{aligned} P_{tot} = -HLC((1 - C_T)T_{ex} - T_0 + C_T T_{fix}) \\ - \tfrac{A_{sol}}{\eta_{HS}} I_{sol} - \tfrac{\eta_B}{\eta_{HS}} P_B + P_B \end{aligned} \tag{3.30}$$

Finally, the total power demand accounting for non-heating periods can be calculated using eq. 3.20, substituting eq. 3.30 for the power during the heating regime.

3.12 Summary

This chapter has presented theoretical models for internal temperatures, conductive, radiative, and ventilation heat transfer, internal gains for dwellings. Linking functions were given to define the relation between heating energy demand and metered energy demand. A formulation of the heat

loss coefficient as the derivative of the total energy demand with respect to ΔT was given. This model enables the operationalisation of thermal performance using readily observable dwelling specific data - the total power demand, external temperature, and solar irradiance. The next chapter will present a method for reliably inferring model parameters from this data.

Chapter 4

Method

4.1 Introduction

This chapter describes the method developed in this research to estimate the dwelling thermal performance from occupied dwelling smart meter data (electricity and gas). First, a definition of the concept of thermal performance used in this research and the manner chosen to operationalise this concept is given, using a suitable thermal model. Subsequently the chapter describes the method of inferring parameters of this model in a robust way, with particular focus on the HLC using a steady-state approximation. Assumptions and limitations are discussed.

Section 4.2 gives an overview of the development of the Deconstruct model. It explains the 'natural experiment' approach and methods to infer the model parameters such as heat loss coefficient. Section 4.4 presents methods for data quality filters, and sampling strategies for selecting the appropriate weather conditions following the natural experiment concept. It further defines the method for estimating HLC using low solar gain conditions and a method to calculate base-load power. Section 4.5 explains the assumptions and the limitations of the Deconstruct model in terms of the simplifications made for the empirical energy model as well as simplifying assumptions for the weather dependencies.

4.2 The Deconstruct method

This research introduces the 'Deconstruct' method, which is a name given to the combination of:

- A grey-box physical model of building thermal performance described in Chapter 3.
- A simplified model of dwelling internal temperature described in Chapter 3.
- A data sampling method.

- A model-fitting algorithm to infer thermal and temperature model parameters.

This model applies to individual dwellings and is suitable to inferring thermal parameters of large numbers of dwellings. In Deconstruct the concept of thermal performance is operationalised using the relation between metered power demand and weather conditions. In general terms, from theory and experiment, domestic power demand for heating is known to be affected by:

1. Building parameters, particularly size, insulation, type of dwelling (detached, semi-detached, flat, etc...)
2. The heating system efficiency
3. Occupant interactions - control systems, strategies, and preferences for hours of heating and internal temperature.

Energy labels like the EPC as well as the PRISM and PTG approaches aim to separate occupant inputs from the building and heating system efficiency, while co-heating tests eliminate occupants altogether by testing empty dwellings. Deconstruct aims to infer unseen parameters separating presumed physical variables from occupant-driven ones by using known building physics to describe the relation between weather conditions and power demand.

The method aims to be scalable, in the sense that it may readily be applied to large numbers of dwellings without incurring significant manual effort, costs, or being computationally prohibitive. Minimising the need for manual intervention implies that the approach should be robust with respect to data quality and the effects of confounding factors, such as occupants. This motivates the use of a grey-box model that links measured energy to thermal parameters, avoiding the need for internal dwelling measurements beyond the smart meter.

In Section 3.11 the Heat Loss Coefficient (HLC) was introduced and its relation to the Heat Transfer Coefficient (HTC) discussed. The HLC combines elements 1. and 2. of the heating demand determinants defined above. The grey-box dwelling energy demand in the heating regime was given in eq. 3.30, and is summarised below in eq. 4.1.

$$\begin{aligned} P_{tot} = &-HLC((1-C_T)T_{ex} - T_0 + C_T T_{fix}) \\ &- \tfrac{A_{sol}}{\eta_{HS}} I_{sol} - \tfrac{\eta_B}{\eta_{HS}} P_B + P_B \end{aligned} \tag{4.1}$$

Deconstruct enables the inference of HLC, P_B, T_0, and A_{sol}.This research focuses on the HLC as a measure of thermal performance. This parameter incorporates the transmission, ventilation, and heating system losses of the dwelling, thereby covering key physical determinants of dwelling heating demand.

Parameters C_T, η_B, and η_{HS} must be set before it is possible to infer HLC, P_B, T_0, and A_{sol} from energy meter and weather using a method independent from the energy-weather data. In the absence of internal temperature monitoring, C_T may be set to a common value for all dwellings

under analysis (for example a nationally representative value). An approach to determining a value for C_T is presented in Section 4.4.3.2. The base-load gain parameter η_B can be set to 1 as a first approximation as previously discussed in Section 3.8.1, although this assumption may be challenged (see Section 6.4.1). The heating system efficiency η_{HS} may be set using data or assumptions. Note that Deconstruct makes it possible to infer HLC without setting a value for η_{HS} (see Section 4.4.3.1).

Deconstruct requires at least one year of daily average data including total energy, external temperature, and solar irradiance for a dwelling. An approach for associating weather data with energy data based on site location is presented in Chapter 5, as well as for removing missing and error values. Deconstruct does not require an error-free dataset, however a rule-of-thumb approach may be applied requiring sites to have no more than 50% missing daily values.

The total power is calculated as the sum of metered energy sources, most commonly electricity power P_{elec} and gas power P_{gas}. In order to calculate total energy demand, it may be necessary depending on the data source to convert volumetric gas readings to power using the gas energy content equation given in eq. 3.15. This is further discussed in Section 5.5.3.3.

The base-load P_B is estimated from a subsample of this data selected using the filters described in Section 4.4.2, which are designed to select non-heating days.

Next, the HLC and T_0 parameters are inferred following the approach in Section 4.4.3.1 using as input a low solar gain subsample of the data selected using the filtering approach described in Section 4.4.3 and the base-load power P_B. HLC may be estimated from a sample of dwelling energy data under conditions of low solar gain, enabling it to be estimated independently of A_{sol}.

Following the calculation of the above parameters, A_{sol} may be inferred using HLC,P_B and T_0, and set values for other parameters as discussed above. However, this was considered outside of the scope of this research.

A summary of the model inputs and outputs and the source of each parameter (input data or model inference). can be found in Table 4.1, while a summary of the steps is shown in Figure 4.1.

Table 4.1: Summary of model parameters and source of parameter value.

Symbol	Description	Source
P_{tot}	Total measured dwelling power demand	Sum of electricity and gas smart meter power P_{elec} and P_{gas}
T_{ex}	Ambient external air temperature	Gridded weather data, using site location

Symbol	Description	Source
I_{sol}	Solar irradiance	Gridded weather data, using site location
P_B	Base-load power appliances, lighting hot water, and plug loads	Estimated from high solar gain sample
HLC	Fabric and heating system loss rate hot water, and plug loads	Estimated from low solar gain sample sample
T_{fix}	External reference temperature	Set to 5 $^\circ$C
T_0	Internal temperature when the external temperature equals T_{fix}	Estimated from low solar gain sample
C_T	Covariance of internal and external temperature	Set using national data or T_{in} data
A_{sol}	Effective solar aperture	Estimated from low temperature sample
η_B	Fraction of base-load contribution to dwelling thermal balance	Set to 1 for simplifying assumption
η_{HS}	Heating system efficiency	Set from dwelling metadata, ignored if only HLC is needed

Fit metrics are calculated including Coefficient of Variance of Root-Mean-Square Error (CVRMSE), Normalised Mean Bias Error (NMBE), and HLC parameter standard error and percentage standard error ($100\sigma_{HLC}/HLC$). Cut-offs are applied to these metrics to determine whether the model should be considered to suitable describe a site's energy demand pattern. The cut-offs are set to 30% for CVRMSE and 30% for HLC standard error. This is discussed in depth in Section 6.3.

4.3 Parameter inference approach

In order to estimate dwelling performance, Deconstruct requires an approach to infer the grey-box model parameters from dwelling data. A common approach is to infer all model coefficients simultaneously using a suitable fitting algorithm. For example in co-heating tests HTC and solar gain parameters are estimated using triaxial regression (Bauwens and Roels 2014). For models with more parameters and potentially non-linear terms, curve-fitting algorithms may be used to infer model parameters. However, energy demand measurements from occupied dwellings measure a range of processes and activities within a dwelling. In addition to heating demand, appliances, cooking, hot water, and more are also captured. This results in a much noisier signal from which the model must infer the desired dwelling properties. Additionally, when fitting models to such data, having many parameters in the model can result in increased parameter uncertainties and biases

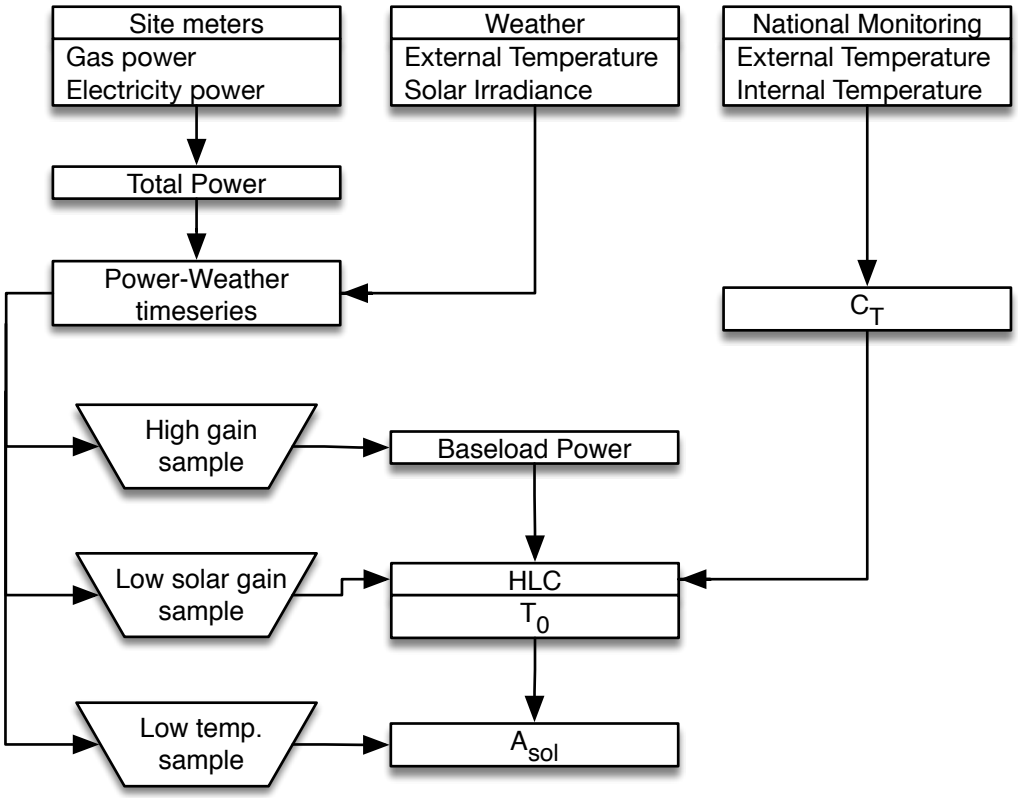

Figure 4.1: Overview of Deconstruct method inputs and steps for parameter estimates

in parameter estimates. For example, Stamp (2015) observed that the inclusion of a solar gain parameter when estimating HTC in a co-heating test could cause a systematic bias in the estimate. Similarly, from this researcher's experience exploring model fitting to smart meter data, it was found that model fit algorithms frequently fail without manual tuning of optimisation algorithm parameters and manual pre/post processing of data.

Instead of attempting to infer all parameters simultaneously, subsets of data are selected according to a 'natural experiment' or 'post-hoc control trial' methodology. A natural experiment is one in which conditions that would be desired for a controlled factor experiment arise naturally within a dataset. To clarify, in order to infer a model parameter, the classic approach would be to take measurements under conditions which minimise unwanted variation.

The normal operation of dwellings is not a controlled experiment - smart meter data is collected without any control over the independent variables. Weather and occupancy conditions vary constantly, and the analyser of the data has no control over these conditions. However it is simple and cost-effective to accumulate data for a dwelling over a long period of time. Given the natural variability of weather conditions during a year, suitable conditions for estimating dwelling thermal parameters will arise 'naturally'. The Deconstruct method identifies data samples that meet conditions suitable for inferring model parameters. Samples of days are selected from daily average data

for a year for a given site and the model fitted to these samples. Where appropriate, the model is simplified when some of the physical processes can be considered negligible (e.g. solar gain terms under low insolation). Using established physics and results from the literature, suitable filters were defined to avoided over-fitting of data by limiting the assumptions made for filtering (Section 4.4). This enables the ingestion of noisy input data incorporating several building operation regimes and selecting from it an optimal subset for the thermal model. The sampling strategy and model fitting approach is described in Section 4.4.

It is important to note that the nature of the Deconstruct data sampling approach precludes the use of dynamic methods. The samples used for the Deconstruct parameter inference are not contiguous, and therefore cannot but used for dynamic models , as these model the dwelling state as a function of (among other things) the preceding time steps i.e. the model state at time step t is a function the state at step t - k. Selecting an arbitrary sub-set of data points to apply the model to is therefore impossible as contiguous data points in the time series may not be dissociated. A steady-state approximation must therefore be used, as the definition of the steady-state limit is that data points are not correlated in time (Rabl 1988). Moreover, the data sampling approach described here is key to achieving robust model fit with minimal manual intervention over large numbers of sites. Observation of dwelling energy data reveals a lot of noise and variation. In this context, to obtain a good fit with a dynamic model requires a) extensive data cleaning and b) incorporating terms in the model to describe the processes generating the data. The first point implies significant labour which would be difficult to automate - achieving sufficient data quality for a steady state model (which by definition will smooth out the most complex transient effects) was already challenging. The second point implies potentially significant increase model complexity and introduces issues around mode under- and over-definition and model identification. It may ultimately be interesting to develop solutions to the problems with applying dynamic models and compare the results, however this is outside of the scope of this research.

Additionally, by not requiring continuous data series, the method is more robust to missing data. Experience in processing dwelling data has shown that such data is frequently subject to data quality issues, such as missing and erroneous data points, unless the utmost care is taken in its collection and archival. Dynamic models require the elimination of error points and filling of gaps in the time series before analysis, for example using interpolation. It may not be possible to fill longer gaps in data, requiring the analysis to deal with discontinuous time series. This would also make the application of dynamic models impossible, and again highlights the advantages of steady-state approaches for this use case.

Daily average data is used throughout this research based on previous findings (Rabl 1988; Chambers, Oreszczyn, and Shipworth 2015). From the literature and extensive exploratory analysis it was found that the underlying daily cycle of demand drivers (day/night cycle and heating patterns) resulted in daily averages being good approximations for steady-state. Figure 4.2 backs this argu-

ment by showing that the model goodness-of-fit as measured by CVRMSE improves rapidly with increasing averaging periods up to 1 day, then does not improve significantly with longer averaging periods. For dwellings with very high thermal mass such as those built to passive-house standards this assumption does not hold - this is a subject for future research. Furthermore, work to compare energy demand between weekdays and weekends found no statistically significant difference (see Chapter 5, Figure 5.15), indicating that no consideration of day-of-week is required.

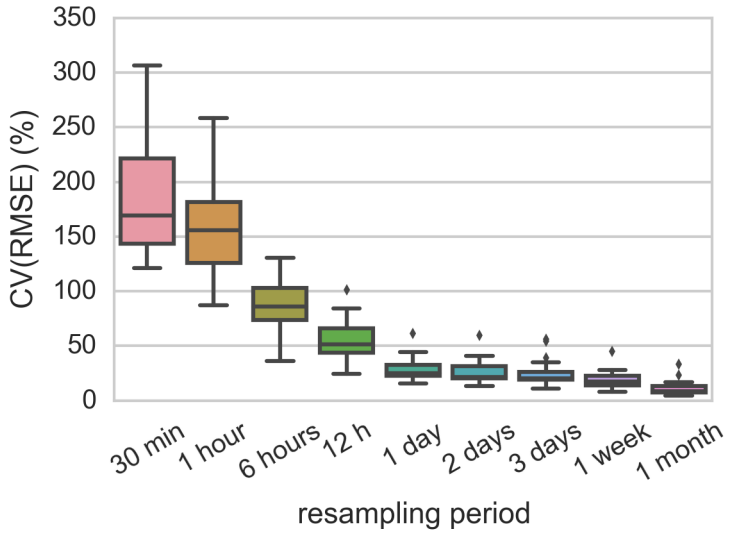

Figure 4.2: Plot of CVRMSE of a linear power-temperature model as a function of resampling frequency demonstrating rapid decrease in residual error with decreasing sampling frequency (longer averaging periods).

CVRMSE (%) is used following ASHRAE recommendations as an overall test for whether the model is representative of data. A cut-off may be defined for CVRMSE above which the model is not considered to be suitably calibrated - we use 30% following ASHRAE recommendations. This enables a first order filtering out of sites which may have bad data or for which the linear model is not appropriate.

Parameter standard errors can be calculated for linear and least-squares fitting problems and gives a first order indication of parameter uncertainty. Further work in Section 6.3 compares this against a more robust MC uncertainty approach and finds it tends to overestimate the uncertainty, however this approach is considerably faster making it useful for general quality control. Figure 6.42 suggests that sites with HLC percent standard error $< 30\%$ should have 'true' (MC inferred) error of $< 20\%$.

4.4 Data sampling for parameter inference

There are two tasks for the data sampling filters defined as part of Deconstruct. The first task is to clean the data and limit the introduction of confounding factors, thus insuring sufficient data

quality for the method to be applied. The second is to select the appropriate weather conditions following the natural experiment concept - finding the weather conditions which a controlled experiment would have selected as a monitoring period. Note that some data quality filters are applied after the natural experiment sample filters.

This section first outlines the data quality filters, then explains the data selection and parameter inference approach for the base-load, HLC, and temperature model coefficients.

4.4.1 Data quality filters

This section describes the data quality filters to improve data namely error filters, outlier removal and absence filters used to improve data.

4.4.1.1 Error filter

Error code columns for gas, electric, and total power were created when the datasets were imported and contain an integer flag for rows with Not a Number (NAN) readings or known errors. All points were filtered where the total error code does not equal 0. Days where the total average energy demand is 0 are considered to be errors. It is more likely that monitoring equipment has failed or that errors have been introduced into the data rather than it being an occupied dwelling consuming no energy at all. Note however that with more reliable data sources a zero daily consumption may represent a genuine measurement - in this case it may be decided not to mark the reading as erroneous.

4.4.1.2 Outlier removal

An outlier removal filter was applied after selecting a sample of daily average data points for the desired weather conditions. Deciding what constitutes an outlier and correct identification of outlier data points is a complex topic (Barnett and Lewis 1978). Iglewicz and Hoaglin (1993) distinguish the three following aspects with regards to outliers:

1. outlier labelling - flag potential outliers for further investigation.
2. outlier accommodation - use robust statistical techniques that will not be unduly affected by outliers.
3. outlier identification - formally test whether observations are outliers.

Outliers in smart meter data arise from measurement issues creating nonsensical data points (such as negative readings) and from unusual conditions within the building causing non-representative energy consumption patterns. The outlier removal filter aims primarily to address the former -

removing data points which are unlikely to have been generated by a normal process and instead are the result of measurement or data handling issues. This avoids subsequent modelling being affected by outliers, particularly in the case where the input data is of poor quality, as there exist many readings far outside the expected range.

Existing literature determined that the modified Z-score approach is suitable for outlier identification and removal (NIST 2013; Iglewicz and Hoaglin 1993; Nascimento et al. 2012). Simple upper limit cutoff filters were found to be inappropriate because dwelling mean power demand distribution is very skew. No single cutoff would be appropriate for all dwellings since it may be much too high for some but much too low for others. The solution is to calculate a cutoff using the individual dwelling power distribution using the median power (the mean would be inappropriate for skew distributions) - this is the basis for the modified Z-score approach.

The Z-score of an observation is defined as

$$Z_i = \frac{Y_i - \bar{Y}}{\sigma}$$

with $\bar{Y}$ denoting the sample mean and σ the sample standard deviation. Z_i is therefore the number of multiples of σ that an observation Y_i is from the mean.

Although it is common practice to use Z-scores to identify possible outliers, this can be misleading (particularly for small sample sizes) due to the fact that the maximum Z-score is at most $(n-1)/\sqrt{n}$. Iglewicz and Hoaglin (1993) recommend using the modified Z-score

$$M_i = \frac{0.6745(Y_i - \tilde{Y})}{MAD}$$

with $\tilde{Y}$ denoting the median and MAD denoting the median absolute deviation, that is the median of the differences between the measured points and the median of the measured points

$$MAD = median(|Y_i - \tilde{Y}|)$$

The authors recommend that modified Z-scores with an absolute value of greater than 3.5 be labeled as potential outliers.

The use of medians rather than means has the advantage of making the filter relatively robust to highly skewed distributions - these are common in energy data where a small number of readings may be hundreds of times larger than typical readings. It is relatively common to apply arbitrary cutoffs to solve this - however since non-outlier dwelling energy demands themselves are a skew distribution, there is no single 'reasonable' cutoff that works across dwelling types and consumption

profiles. One may attempt to work around this by using a multiple of the dwelling's mean or median consumption instead of an absolute cut-off, however this approaches the modified Z-score method in any case without the more rigorous theoretical underpinnings of the latter.

It is important that the outlier filter is applied only after filtering for errors, especially erroneous zero values, because a large mass of these values can affect the median calculation and result in incorrect identification of outliers.

4.4.1.3 Absence filter

The heating model will not apply when a dwelling is not heated. These periods generally coincide with unoccupied periods - therefore an 'absence' filter was developed. Developing a rigorous approach for this is difficult as there is no monitoring of occupant presence. Additionally, making strong assumptions about occupancy is avoided which could result in biased parameter estimates.

The approach to removing unoccupied days for each sites' data was as follows. In a first instance, the daily average energy demand was required to be greater than zero. This was a reasonable requirement for a dwelling in the heating regime. However, only removing points where the energy consumption is zero is not very effective, since a range of appliances use power at all times (notably fridges and freezers). It was therefore further assumed that for a dwelling with a normally operating heating system, heating demand should increase monotonically as external temperature decreased. It follows that in the ideal case, energy demand at a higher external temperature should be lower than at a lower external temperature. Therefore a filter was applied which selected a power cut-off value such that the power demand increases with decreasing internal temperature, using a threshold defined as follows. The power demand P_{Tmax} for the highest temperature data point in the sample was used as a lower cut-off threshold for the power value, removing all points where $P_{tot} < P_{Tmax}$ (Figure 4.3).

As the occupancy state is effectively inferred from the energy demand, dwellings which are occupied but for some reason not heated will appear identical to unoccupied dwellings. From the perspective of inferring the HLC this distinction does not have an impact, since only the change in demand with change in weather conditions matters.

This cut-off has proven to be both simple and effective. The followring approaches were attempted but were not found to be effective:

- Elliptic-envelope anomaly detection (Ng 2013). This approach fits a two dimensional gaussian probability distribution around the data and rejects points that appear too unlikely. This has two problems. Firstly, the Gaussian fit can be heavily influenced by larger groups of 'out of distribution' points. This results in inflated distribution widths which in turn means that all points end up being within the acceptable region.

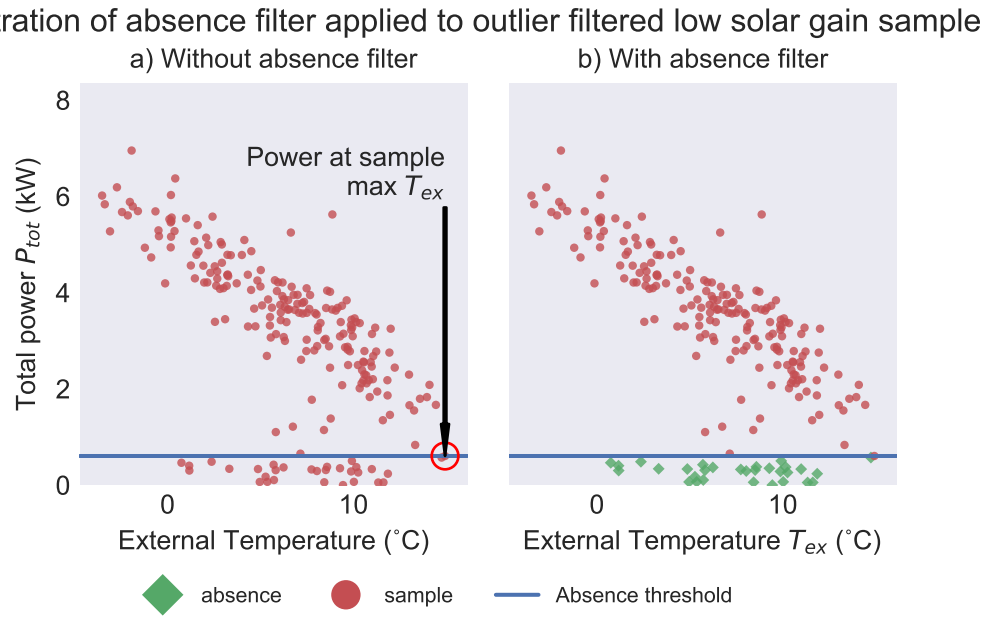

Figure 4.3: Illustration of absence threshold filter applied to a lower-solar gain sample with errors and outliers removed. a) shows the input sample with the power demand P_{Tmax} for the highest temperature data point in the sample highlighted (red circle and annotation). b) demonstrates the effect of the absence filter by showing the resulting output sample and the removed data points.

- Using the mean power during the summer months as the lower cutoff. The disadvantage of this approach was that dwellings had highly variable base-load profiles - in some cases the cutoff was too small to remove unwanted points while in others (especially where base-load variance was high) the cutoff was rather high. Using the upper 95% confidence interval level of the summer mean load instead of the mean displayed similar issues.

4.4.2 Base-load power

Base-load power is defined as the mean power outside of the heating regime, where power is independent of weather conditions and normally distributed about the mean. The requirements for this sample are comparatively flexible, as they only need to guarantee that data comes from outside the heating regime, not capture every data point outside this regime. The approach is to require high solar gains and a minimum temperature cut-off such that heating is not expected to operate during the selected days. The filters applied are, in the following order:

1. Remove missing values and errors (Section 4.4.1.1)
2. High solar gains filter, $I_{sol} > 100W/m^2$
3. High temperature filter, $T_{ex} > 18\,^{\circ}C$
4. Modified z-score outlier removal filter (Section 4.4.1.2)

Base-load power P_B is the mean of the resulting sample.

4.4.3 Low solar gain conditions

A data sample selection with low solar gain is suitable for estimating HLC and internal temperature model parameters. Limiting solar gains greatly simplifies the model since solar gains introduce complex time-of-day, time-of-year, and building geometry factors into the energy model. Since dwellings have a given orientation and exposure of their different surfaces, the amount of energy they receive from the sun will depend on how they are exposed over the course of the day, which in turn is affected by the path of the sun through the sky which changes over the year. Minimising solar gains makes these effects negligible. Maintaining near-constant conditions also mitigates the issue of the potential covariance of occupant-building interactions with weather. Specific concerns were raised about the likelihood of occupants systematically opening windows in autumn and spring on bright but cold days (Oreszczyn 2014) - this should not be the case under low solar gain and low temperatures.

Together with sampling days with low solar gain, the filters required to clean the data of errors and outliers, as well as to account for edge cases and ensure only days where the building is occupied are included, were applied in the following order:

1. Error filter. This filter removes points with missing data, see Section 4.4.1.1.
2. Low solar irradiance filter. This filter selects days where the mean solar irradiance is below a given cut-off value of 0.05kW/m2. Solar gains can contribute significantly to a dwelling's energy balance and bias the result. This filter is based on work by R. Everett, Horton, and Doggart (1985a); Siviour (1981); Stamp (2015) and is particularly effective in the UK because the weather is often cloudy.
3. Outlier removal filter, according to Section 4.4.1.2.
4. Temperature backstop filter. This filter eliminates days where T_{ex} is greater than a cut-off temperature. This is a 'backstop' filter as it aims to eliminate unusual conditions that are not already removed by the primary filters - there are generally very few days with high air temperatures but no sun. The upper temperature cutoff is set to 15 $^\circ$C following work by Ruch and Claridge (1992), Summerfield, Oreszczyn, and Hong (2013), Hamilton et al. (2016a), Oreszczyn et al. (2006) which showed a statistically better fit to consumption data below about 15°C, while also limiting the overall impact on gradient estimates (Figure 6.10). This also increases the minimal ΔT between inside and outside, resulting in a corresponding reduction in the fractional uncertainty in ΔT.
5. Absence filter, see Section 4.4.1.3.

The sequential application of these filters is illustrated using a single site in Figure 4.4. As can be seen from the graph, the solar gains filter has the largest effect, while the other filters serve to remove a relatively small number of unusual points.

It is critical that the outlier filter be applied only after the desired subset of data has been selected

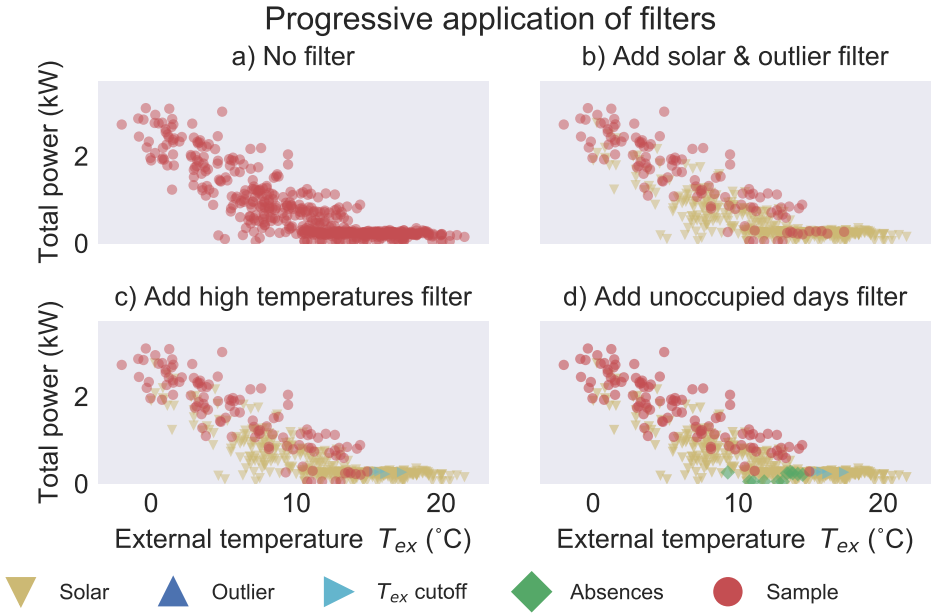

Figure 4.4: Illustration of the effects of adding each filter in turn to the data. a) Data with no filtering. b) Sample after application of the low solar gain filter and outlier removal. c) Filter for external temperatures above a 15 ˚ C cutoff, a small number of points are removed. d) Filter for unoccupied days, which removes a number of low consumption points.

through the low solar gain filter, because the full dataset includes several operating regimes, while the outlier filter does not make assumptions about the underlying model. Before applying the solar gains filter, the energy data in effect contains two distinct power distributions corresponding to a non-heating regime and a heating regime. This results in the median value of the power points no longer having a clear meaning. In this condition, the data from one operating regime may appear as an outlier relative to the data for another regime. For example, power demand during the heating period looks like outlying data relative to the energy demand in summer, as illustrated in Figure 4.5. In this illustration we see clearly that the outlier filter incorrectly identifies a large fraction of the winter demand as outliers. When the solar filter is applied first, no points are identified as outliers, because the energy demand distribution corresponds to a single operation mode.

Using the low solar gain condition, eq. 4.1 may be simplified to eq. 4.2:

$$P_{tot} = -HLC(1 - C_T)T_{ex} + HLC\ T_0 - HLC\ C_T T_{fix} + P_B(1 - \tfrac{\eta_B}{\eta_{HS}}) \tag{4.2}$$

4.4.3.1 Heat Loss Coefficient

The HLC can be estimated using the gradient of eq. 4.2 with respect to external temperature:

$$\frac{dP_{tot}}{dT_{ex}} = -HLC(1 - C_T) \tag{4.3}$$

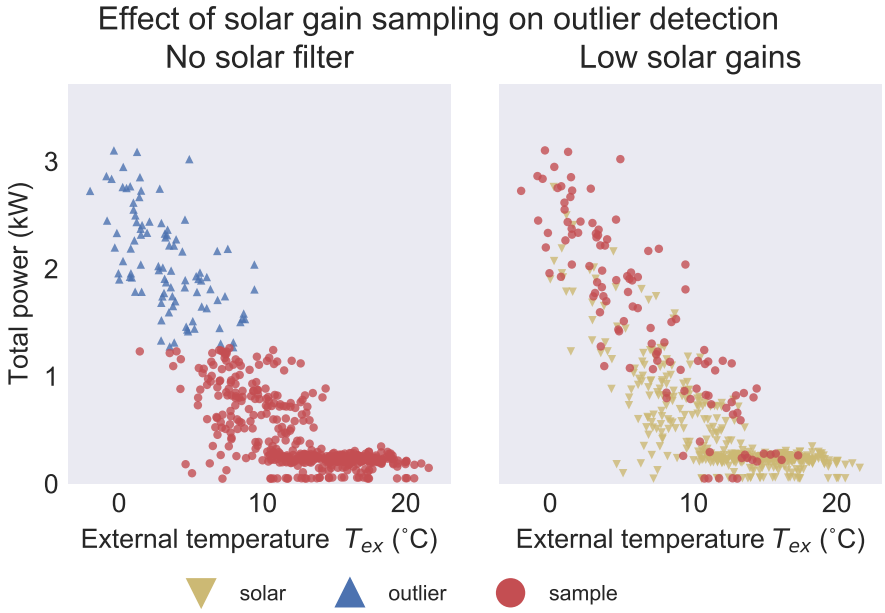

Figure 4.5: Illustration of applying an outlier filter to all input data. Application of outlier filter to unfiltered dataset (left) removes valid high consumption points from cold days. Removal of high solar gain days first (right) reduces the dataset and results in no points being identified as outliers.

which also demonstrates that the HLC cannot be estimated separately from C_T without additional data. This will be addressed in the following section. Note the relation between eq. 4.3 and the Power-Temperature Gradient (PTG) (Summerfield, Oreszczyn, et al. 2015a), where $\frac{dP_{tot}}{dT_{ex}} = -PTG$ but with a different sampling approach where only a temperature cutoff was applied. This approach will take into account energy gains from metered primary and secondary heating and base-loads, but will not be able to account for non-metered heat gains, as further discussed in Section 4.5.

4.4.3.2 Internal temperature model parameters

The internal reference temperature T_0 (internal temperature when external temperature is 5 ˚ C) is related eq. 4.2 at absolute zero (as a reminder, calculations are performed in Kelvin although results are expressed in ˚ C) $T_{ex} = 0$ ˚ K. T_0 is given a function of the power at absolute zero, $P_{tot,0}$, and the building thermal parameters in eq. 4.4.

$$T_0 = \frac{1}{HLC}(P_{tot,0} - P_B(1 - \frac{\eta_B}{\eta_{HS}})) + C_T T_{fix} \tag{4.4}$$

HLC is obtained as described in the previous section while $P_{tot,0}$ is obtained from the y-intercept of a linear regression, while P_B is estimated as described previously in Section 4.4.2. The temperature covariance C_T, cannot be estimated using dwelling energy data due to it being co-linear with HLC (eq. 4.3).

We must therefore determine C_T by other means. When internal temperature is available, both C_T and the mean internal temperature at 5 °C T_0 can be estimated from monitored internal temperatures. Doing this may appear to defeat the purpose of modelling internal temperatures - the measured temperatures could instead be used directly. However, using this temperature model makes it simple to fill gaps in measured data, as well as reduce the effect of noise and outliers. Dependence on measured internal temperatures would still greatly limit the application of the approach.

In the case that internal temperature measurements are not available, it may be preferred to simplify the problem by setting $C_T = 0$, i.e. assuming that the covariance is not significant. This assumption is likely to be invalid however, due to the internal temperature change being driven by the length of heating periods due to the fact that during cooling periods T_{in} trends towards the external temperature.

The suggested approach to obtaining C_T for dwellings without T_{in} measurements is to use a separate dataset to estimate a nationally representative mean C_T for UK dwellings by using measured internal temperatures from the Energy Follow Up Survey (EFUS) dataset (described in Chapter 5), which studied a nationally representative cross section of dwellings. Once C_T has been estimated, the T_0 may be inferred when fitting the power demand model to data.

A comparison of these approaches is given in Chapter 6.

The η_B and η_{HS} parameters remain unknown. T_0 depends linearly on the base-load gain factor η_B ($\frac{dT_0}{d\eta_B} = \frac{P_B}{HLC}$). A first order approximation is to assume that all base-load energy use is eventually dissipated as heat into the dwelling fabric, based on the second law of thermodynamics. This assumption is reasonable in the case of lighting and appliances but may be less so in the case of water heating as hot water may be lost to drains without heating the dwelling - this would depend on the amount of hot water consumption relative to base-load power.

As previously discussed, η_{HS} cannot be inferred independently from HTC without additional data, and unlike η_B there is no theoretical basis for supposing the value to be 1, outside of the case of electric resistance heating. The effect of the choice of η_{HS} on T_0 is investigated in Section 6.4.2 for a set of dwellings with known boiler efficiency ratings.

4.5 Assumptions and limitations

4.5.1 Model simplifications

The thermal model of the dwelling is by necessity very simple, using a single thermal zone. Thermal bridging or transfer to the ground was not considered, because suitable data is not expected to be collected for most dwellings. Deconstruct may only provide estimates of lumped dwelling

parameters. These alone are not sufficient to predict the effect of specific retrofits (e.g. installing double-glazed windows). Without additional dwelling information, it is difficult to do better. However, by referring to the component thermal balance equations, it would be possible to factor out various parameters given additional information about the dwelling.

Unmetered thermal gains apart from solar gains are difficult to assess given the lack of data. Metabolic gains (from occupants) are assumed, like occupancy pattern, to be stable and non-correlated. Gains from occupants are relatively small - SAP gives typical gains of 60W per occupant (BRE 2011). At a UK average of 2.3 people per household or 138W added when occupants are present, this is small compared with heating and solar gains of several thousand watts. When estimating the HLC we therefore assume negligible impact from these gains.

Secondary heating gains will be included as a matter of course if the heating systems use a metered energy source such as electricity or gas. Unmetered heat inputs from additional heating fuels such as wood fires are not possible to measure, in fact it isn't possible to establish their existence without an on-site visit or further information supplied by residents.

Occupancy profiles are not modelled, beyond the filter which removes days where the building appears to be unoccupied. Occupancy should be a critical driver of energy demand, while modelling of occupancy is a large topic in itself. However, the source of the Deconstruct method is the observation (going back to the PRISM method in 1986) that energy demand follows a reliable pattern across days without attempting to take detailed occupancy profiles into account. It appears that mean occupancy is a stable quantity, although without detailed occupancy data this can only be inferred by the stability of dwelling power demand.

The Deconstruct method relies on a sufficiently long monitoring period for the desired conditions to arise in the dataset - at least a full winter of data for HLC estimation and a full year for other parameters. This drawback is offset by the simplicity and low marginal cost of using smart meter data - unlike other approaches, collecting this relatively long period data does not incur significant additional cost or effort. On the contrary, it presents an opportunity for governments and utility companies to derive additional value from data that they are already collecting and archiving. In the UK, smart meters already store by default data for 13 months in their internal memory (Beard 2015).

4.5.2 Weather variable assumptions

From theory and empirical evidence it was decided not to include the following weather parameters in the final model.

Wind speed might be expected to have an impact on energy demand, based on the theory in Appendix B and measurement work in R. Everett, Horton, and Doggart (1985b). However, work

presented in Section 6.1.2 demonstrates that there is empirically no significant dependence of power demand on wind speed either overall or for the low solar gain sample. Wind direction is non-trivial due to the relation between 'global' wind direction and local wind direction in urban areas. The influence of wind would also not be a monotonic function of direction, further complicating modelling. Given the lack of dependence on wind speed alone, it was deemed that developing a wind direction model was unnecessary within the scope of this research.

Data analysis presented in Section 6.1.3 suggested no significant correlation between low solar days and rain which might have had an impact on energy demand. Additionally, there is a limit to how low the low solar gain filter may be placed without removing too many data points - the chosen cutoff was based on previous work in this field.

Weather dependence of appliance energy use is not considered. Correlation of total load with temperature in summer months presented in Section 6.1.1 suggests that there is little correlation between appliance load and weather. Lighting demand would also be expected to change seasonally, however lighting energy should simply contribute to the dwelling heat balance. Nonetheless, the low solar gain sample used to infer HLC reduces the expected change in lighting demand as illumination conditions should remain constant. It is possible that if smart appliances become prevalent in the future, and these appliances are used for demand response, that their consumption may become weather-dependant as variation in energy generation from renewable sources such as solar and wind trigger demand response actions from the appliances. However, such developments have yet to come to pass.

4.6 Summary

This chapter has provided a working definition of the concept of thermal performance and has operationalised this concept as the parameters of a model relating measured domestic energy use and weather conditions. It has introduced 'Deconstruct', a method for estimating the parameters of whole-building energy balance equations from occupied dwelling data using a 'natural experiment' approach. This approach aimed to take advantage of the large amounts of data that may be gathered with relative ease using smart meters by discovering subsets within that data corresponding to conditions equivalent to those that would be chosen for a control trial. A simplified model for the thermal energy balance was defined, highlighting the Heat Loss Coefficient (HLC) as a key indicator of whole dwelling efficiency, measuring a combination of the dwelling fabric performance and the heating system efficiency (eq. 3.26). The internal temperature model presented in Chapter 3 (eq. 3.4) was integrated into the thermal model. A data sampling strategy was developed to enable model parameters to be inferred. A low solar gain sample was proposed to estimate HLC and internal temperatures while eliminating bias and complexity caused by solar energy gains. A simple sample to estimate base-load power using non-heating season data was also presented. Data

quality filters required to address errors, outliers, and unoccupied dwellings for all of these samples were also described. This chapter also included a discussion of steady state approximations, as well as of known assumptions and limitations in the method.

Chapter 5

Datasets

5.1 Introduction

This chapter describes the datasets used in developing the 'Deconstruct' method and model described in this research. First, a brief overview is given of the software framework developed to import, store, analyse, and visualise these datasets in Section 5.2. This aimed to facilitate reproducible research.

Geographical Information System (GIS) data on postal and administrative regions, gas distribution zones and weather reference grids is presented Section 5.3. This information was needed notably to link weather datasets with dwelling monitoring data. Section 5.4 discusses the weather data sources and presents the two weather gridded datasets used for this research - the Climate Forecast System Reanalysis (CFSR) from the National Centers for Environmental Prediction (NCEP), and an extract of weather data from UK MetOffice. Two weather datasets were needed in order to cover the full time range of the dwelling data, as well as to provide the maximum spatial resolution.

As discussed in Section 4.4.3.2, the temperature covariance model parameter C_T may not be inferred without internal temperature measurements. Instead, a nationally representative average may be estimated. For this purpose, the Energy Follow Up Survey (EFUS) data was used, which is described in this chapter under Section 5.5.1.

Empirical dwelling data was needed to develop and evaluate the 'Deconstruct' method and model described in Chapter 4. At a minimum dwelling gas and electric smart meter data is needed, with location information in order to link weather data. To test the method, enough metadata must be available for a dwelling to provide meaningful points of comparison. An ideal dwelling dataset would include extensive monitoring of parameters including gas and electricity smart meter readings, internal temperatures, monitoring of boiler operation and efficiency, RdSAP assessment input data, and on-site weather monitoring. It would also have enough sites to allow a portion to be used

for method development and a portion to be reserved for cross validation. Unfortunately no such idea dataset exists. To address this problem, four smart meter datasets were obtained - the Energy Demand Research Project, Electricité de France (EDF) subset of EDRP (EDRP-EDF), Solid Wall Insulation (SWI) field trial, and the Pennyland Project. Each offering a different mix of benefits and drawbacks in terms of the number of dwellings, the monitored parameters, and the amount of dwelling metadata. Section 5.5.2 discusses these in detail, with an overview of the dataset regional and temporal coverage, as well as the availability of relevant metadata, given in Table 5.4. Briefly, the EDRP-EDF dataset was used most extensively for the method development and evaluation given its relatively large size (~1800 sites), accurate location information, and metadata concerning dwelling age, type, and size which enabled comparisons of broad dwelling categories. This was supplemented using the SWI dataset which included internal temperature measurements, although it was much smaller and suffered from data quality issues (see Section 5.5.3.3). The EDRP and Pennyland Project datasets were used for analysis in the Results Chapter 7.

The weather, power, and where relevant internal temperature data were combined using the GIS datasets to perform a geospatial linking between the different inputs and produce power-weather or internal temperature-weather time series. This process is outlined in Section 5.6.

5.2 Software framework

Developing a robust system to manage data is crucial part of the research in order to:

- Facilitate extraction and cleaning of input data into a consistent and readily accessible format.
- Automate loading, table linking, and data query steps.
- Enable process automation and reproducibility, such that the entire data analysis could be re-run from the original data files to the final outputs.
- Be self-documenting through the use of structured and self explanatory code.

This facilitated collaboration and dissemination of datasets and results. The software tools used included:

- The Python programming language (Python Software Foundation 2015), "Jupyter Notebooks" for interactive programming and analysis containing live code, equations, visualisations and explanatory text were used extensively (Perez and Granger 2007). Scientific computing modules used included the following references:
  - NumPy (efficient numerical computation) (Walt, Colbert, and Varoquaux 2011)
  - SciPy (scientific numerical functions) (Jones et al. 2015)
  - Pandas (tabular data processing) (McKinney 2010)
  - Xarray (N-dimensional labelled data structures) (Hoyer and Hamman 2016)

  - Matplotlib (plotting) (Hunter 2007)
  - Cartopy (map plotting) (Met Office 2016)
  - PVLIB (solar irradiance calculations) (Andrews et al. 2014)

- The PostgreSQL database with PostGIS extension. PostgreSQL is a widely used high performance SQL relational database with good support for time series and compact high performance indexes. PostGIS adds support for geographic objects, allowing geospatial queries to run in SQL.

- The Network Common Data Form (NetCDF) file format and associated software modules. NetCDF is a set of software libraries and self-describing, machine-independent data formats that support the creation, access, and sharing of n-dimensional array scientific data (UCAR 2016) and is widely used in meteorology.

An optimised software framework was developed (Figure 5.1) to provide a convenient interface to input datasets and for building process workflows. Documentation was written to facilitate core reuse and reproducibility, in line with emerging recommendations on open scientific software (Software Sustainability Institute 2016).

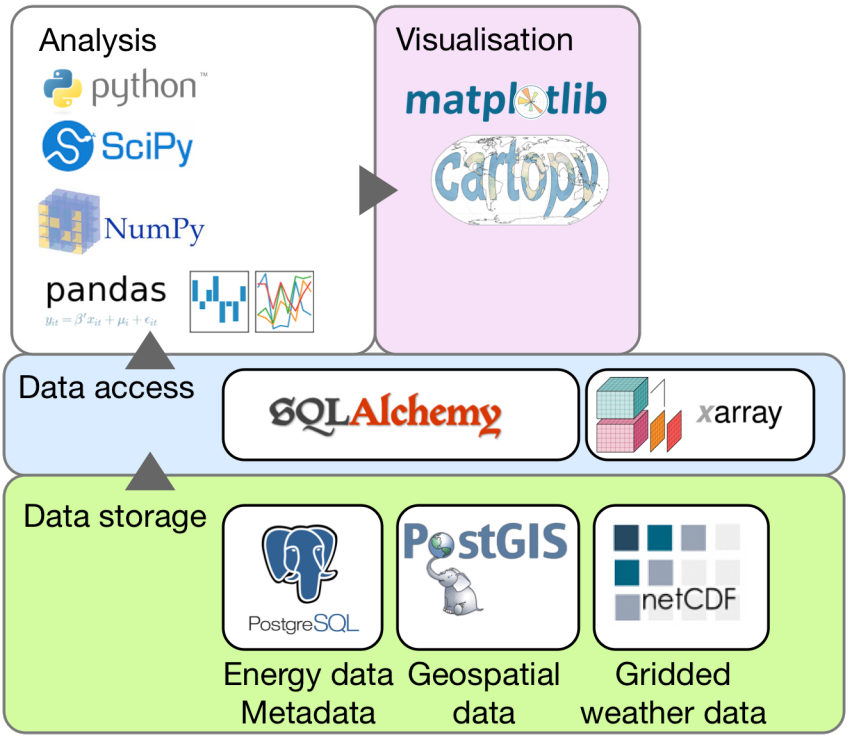

Figure 5.1: Overview of software stack with data storage layer, data access layer, and analysis and visualisation layer

5.3 GIS data

Geographical Information System (GIS) data was needed to in order to link dwelling location metadata with geographical features, fetch corresponding time series, and produce map visualisations. GIS data maps values (numerical or qualitative) to positions (with longitude, latitude coordinates) or areas (list of points defining the polygon outline) on the globe. To ensure compatibility between different GIS software, the World Geodetic System (WGS) 1984/EPSG:4326 coordinate standard was used which defines points on the earth's surface as longitude/latitude angles in degrees East and North respectively.

GIS data is commonly provided in the industry standard 'shapefile' interchange format. Each shapefile generally contains one geographical feature (e.g. a set of country borders) (ArcGIS 2016). These files were imported as PostGIS tables for ease of integration into processing.

5.3.1 Postal and administrative regions

Dwelling metadata contained a full or partial postcode, or an administrative or statistical region code to locate the dwelling. Files are publicly available defining the geographic location of postcodes (ONS 2016), as well as the region polygons describing postcode areas and districts (OpenDoorLogistics 2016), and Nomenclature of Units for Territorial Statistics (NUTS) regions (Eurostat 2016) and Local Authority Districts (LAD) (Ordnance Survey 2016). The coverage and relative size of the geographical regions is shown in Figure 5.2.

In the UK postcodes consist of two parts separated by a space. The first part is the 'district' code. The first one or two letters of district is the area code, which combines multiple districts. For example, in the postcode PO1 3AX, PO is the postcode area of Portsmouth, PO1 is a postcode district within the postcode area of Portsmouth. Knowledge of the allowed format of UK postcodes was important for validating partial postcodes provided with some datasets and resolving ambiguous cases.

5.3.2 Gas distribution zones and time series

The UK National Grid distributes gas through it's network, which is divided into a 13 Lower Distribution Zones (LDZ), the boundaries of which were obtained from official shapefiles. These zones are useful because of their association with officially measured gas calorific values. The heat content of natural gas in the distribution grid constantly varies within an allowed range due to small variations in the composition of the gas mix. Since domestic gas consumption is generally metered in terms of volume, the energy content or Calorific Value (CV) of the gas is needed to calculate the energy used. The National Grid provide time series of gas CV which can be associated

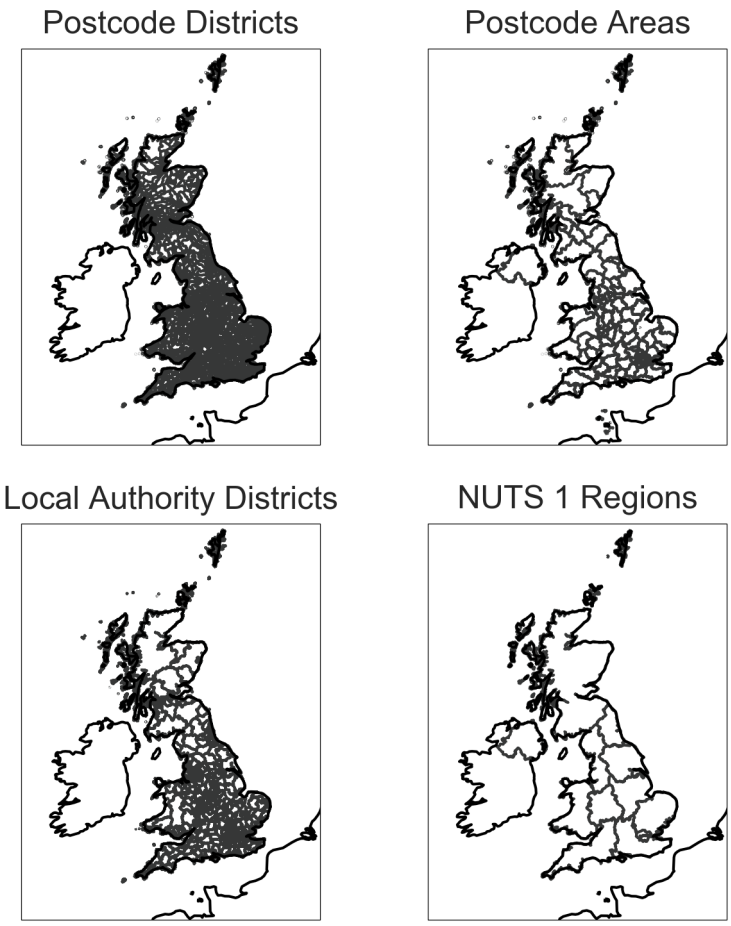

Figure 5.2: Illustration of the regions corresponding to postcode districts and areas, local authority districts, and nuts1 census regions

with the corresponding LDZ. Tables of CV time series were uploaded to the database alongside the LDZ shapefile (NationalGrid 2015). The mean CV was 39.4 MJ/m^3 over all time with a standard deviation of 0.4. For the daily values, the median difference between the daily minimum and maximum recorded CV across the network was 1.2 MJ/m^3, or approximately 3%.

It is important to note that gas energy content by unit volume is affected by temperature. National Grid does not make any particular allowance for this within their network. This is reasonable as most gas pipelines are underground and therefore do not experience large temperature swings. However, changes in temperature and pressure at the meter point may induce uncertainties of the order of 3% (see Section 6.3).

5.3.3 Global weather reference grids

Weather data was stored as gridded time series (a labelled 3-dimensional array) in a netCDF binary file. To enable cross-referencing GIS data with weather, point grids were generated based on the longitude-latitude coordinates from the weather files which matched the weather grid points.

Rectangular areas were generated centred on each point with sizes defined by the grid point spacing in the longitude and latitude directions (Figure 5.3). The area polygons were necessary in order to enable queries looking for the overlap between a search area (such as a postcode district) and weather grid cells. The points and polygons were uploaded to PostGIS as geo-referenced tables.

Note that gridded weather data represents zero dimensional point samples of weather (pixel data). This approach effectively assigns the values of the weather point to an area surrounding it, with no attempt to interpolate these values. This was found to be the most practical and performant solution.

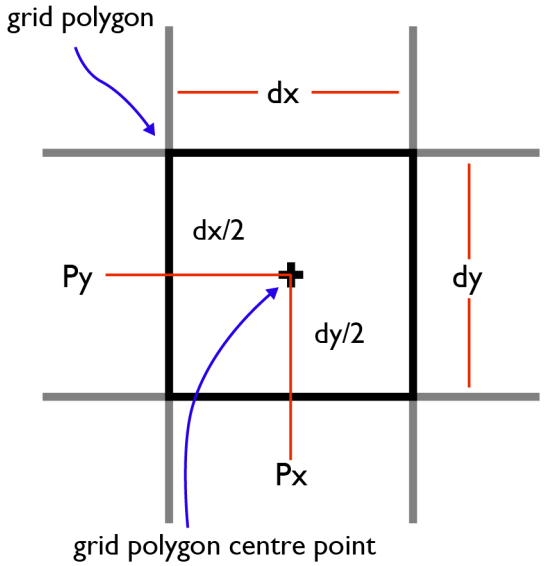

Figure 5.3: Illustration of construction of point and polygon grids for cross referencing weather data. The position (Px, Py) of the polygon centre acts as an index label for the whole polygon

5.4 Weather data

5.4.1 Gridded Numerical Weather Prediction data

The historical weather data used in this research was drawn from Numerical Weather Prediction (NWP) models, on the recommendation of the MetOffice. The alternative would be to use data from the weather station nearest to each site. Weather station data has the following drawbacks compared to NWP data:

- The available variables are limited to those measured at a given site. Some useful variables, such as solar irradiance, are not measured at many weather stations. NWP guarantees every variable of the model is available at every location.

- Weather stations may be relatively distant from a given site. One solution would be to

interpolate between two or more stations - however this is in a sense what NWP models do, except using far more advanced models and a large range of observational data.

- Weather stations are generally placed in locations with few confounding factors such as buildings and urban areas. They are therefore not necessarily representative of conditions inside built-up areas. NWP models include information about ground surface properties and can therefore account to an extent for built-up areas.

- Readings from weather stations generally give daily values (such as temperature high, low, and mean) while NWP models are generally hourly.

- Weather stations are occasionally moved, new ones are introduced and old ones shutdown, which could make data retrieval more complicated (MetOffice 2015). NWP model data is on a uniform grid, greatly simplifying location-specific data lookup.

The principal drawbacks of NWP data are:

- Modelled data will never be as good as on-site monitoring. Weather conditions may vary significantly at small scales, for example certain regions may enjoy microclimates while dense urban areas are subject to a range of specific phenomena that affect local weather not modelled by NWP. This is particularly true of wind speed and direction which can be heavily influenced by the immediate environment of a building (vegetation shelter, configuration of nearby buildings, etc.).

- NWP models are limited in time coverage. MetOffice models in particular are updated every few years, creating discontinuities. Combining time series from different model iterations is not recommended. Projects to develop long term, consistent models by re-processing historical data are also limited by the availability of suitable satellite data and can therefore generally not reach back beyond 1979, while weather station data may go back over 100 years.

- The density of the NWP grid will dictate how location specific the weather data can be.

Managing gridded weather data can be challenging due to its size - one UK-wide grid per hour for several years of data. Storage as NetCDF files is standard in the meteorology community and gives access to a number of optimisations, such as transparent data compression. The Xarray Python module was used to accelerate and parallelise loading and subsetting of 3-d gridded weather NetCDF files, code was contributed to this module as part of this research (Hoyer and Hamman 2016).

Checking errors is also challenging with large files - it becomes unfeasible to check every value in the dataset. For this research, it is assumed that the weather data providers (MetOffice and NCEP) provide error free datasets.

5.4.2 Weather Datasets

Two weather gridded datasets from numerical weather models were used for this research from the National Centers for Environmental Prediction (NCEP) Climate Forecast System Reanalysis (CFSR) and from the UK MetOffice.

Each had specific advantages and disadvantages; ultimately both datasets were needed to cover all use cases in this research. This becomes apparent when cross referencing the weather dataset properties in Table 5.1 with the energy meter dataset properties in Table 5.4 in Section 5.5.2. The CFSR has a larger temporal and spatial range than the MetOffice data. This was necessary in order to obtain weather data corresponding to the Pennyland project time range (1980-1983). However it has a lower spatial resolution than the MetOffice (~0.312 ° vs ~0.036 °), this additional resolution was important for recent datasets with medium and high precision location data - EDRP-EDF and SWI. Additionally, the discontinuity of CFSR in 2010 was inconvenient with respect to the datasets that crossed that time range - EDRP, EDRP-EDF, and SWI.

The lower resolution of CFSR meant that data files were smaller, facilitating experimentation. The access conditions and process via the University Corporation for Atmospheric Research (UCAR) for the CFSR were considerably simpler than the MetOffice, with the latter requiring lengthy in-person coordination with the MetOffice as well as considerable costs. Therefore much of the exploration and development work was performed using CFSR with final result produced using the MetOffice data.

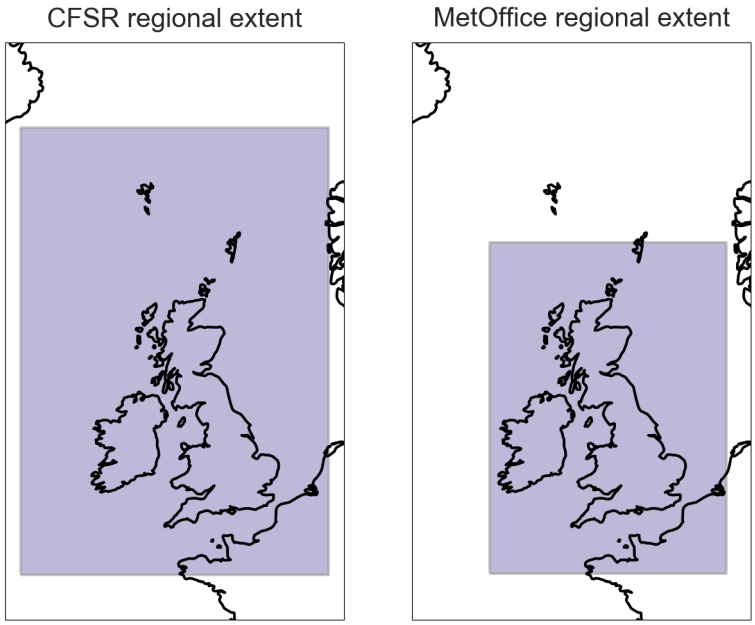

Figure 5.4: Illustration of -15 ° to 5 ° longitude, 48 ° to 64 ° latitude selection region for CFSR data grid and the -10.96 ° to 4.39 ° longitude, 48.05 ° to 60.52 ° latitude region for MetOffice data grid.

Table 5.1: Comparison between properties of the CFSR and MetOffice datasets.

Property	CFSR	MetOffice
Max time range	1979-2010	2008-present
Grid size	~0.312 ° / ~22km	~0.036 ° / ~4km
Grid range longitude	-15 ° to 5 °	-10.96 ° to 4.39 °
Grid range latitude	48 ° to 64 °	48.05 ° to 60.52 °
Model coverage	Global	UK
Data costs	Free for academic use	Free data, cost for handling and delivery
Data access	UCAR archive	Private file transfer
Data format	NetCDF3	Custom binary
Data integrity	Guaranteed continuity	Some gaps

The variables drawn from each dataset are shown in Table 5.2. For CFSR, certain variables were only available on a coarser 0.5 ° grid compared to the maximum resolution 0.312 ° grid. Some of the variables obtained for CFSR were not available for the MetOffice. Figure 5.4 illustrates the geographic range of the datasets, these are very similar.

Table 5.2: Variables obtained from the CFSR (fine 0.312 ° and coarse 0.5 ° grid) and MetOffice datasets.

Variable	CFSR 0.312 °	CFSR 0.5 °	MetOffice
Air temperature at 2m altitude	Yes	Yes	Yes
Relative humidity at 2m altitude	No	Yes	Yes
Dew point temperature at 2m altitude	No	Yes	No
Shortwave horizontal irradiance at ground level	Yes	Yes	Yes
Longwave horizontal irradiance at ground level	Yes	Yes	Yes
Cloud cover from ground level	Yes	Yes	No
Air pressure at ground level	Yes	Yes	No
Wind velocity at 10m altitude	Yes	Yes	Speed only
Precipitation	No	Yes	Yes

A full description of the CFSR project can be found in Saha et al. (2010). The MetOffice data was sourced from the UK variable resolution (UKV) model outputs at a grid resolution of approximately 4km for 2008-2012 over an area covering the British isles. The UKV model is the UK national forecasting model whose outputs are continuously harmonised with observational data (MetOffice 2016).

5.4.2.1 Wind speed

The CFSR weather dataset provided wind velocity as U and V components, where the U direction denotes wind blowing towards the east from the west, and V direction towards the north from the south. Wind speed is therefore:

$$s_{wind} = \sqrt{U_{wind}^2 + V_{wind}^2}$$

The influence of wind speed was investigated in Section 6.1.

5.4.2.2 MetOffice weather data processing

The MetOffice data required a number of processing steps in order to be compatible with the CFSR dataset. The data was provided as MetOffice format '.pp' binary files. Fortran code was provided to read these; this code was ported to Python for compatibility with the rest of the data workflow. The raw files were archived as NetCDF4 files, one file per variable per year, with compression enabled.

The gridded data was provided on the MetOffice 'rotated pole' coordinate system. This coordinate system rotates the classical longitude/latitude system by shifting the polar axis to pass through longitude=177.5 °, latitude=37.5 ° such that the 'equator' and the zero-th meridian intersect in UK. This makes size in km of the grid squares defined by an angle step roughly constant, which would not be the case in standard coordinates due to the UK's northerly location (Figure 5.5). This coordinate system is incompatible with the WGS84 standard. Fortran code was provided by the MetOffice to transform the rotated equatorial coordinates to longitude/latitude coordinates; this code was ported to Python and used to transform the coordinates.

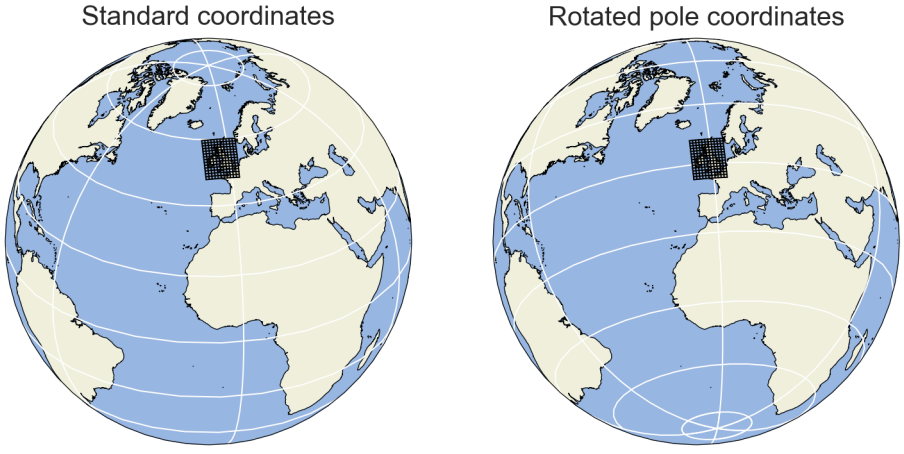

Figure 5.5: Globes showing the standard longitude/latitude gridlines (left) compared to the MetOffice rotated pole coordinate system gridlines (right).

After converting to standard coordinates the grid lines are no longer straight and perpendicular, as seen in Figure 5.6. A new weather dataset was generated by resampling the rotated grid onto a standard longitude/latitude grid with a larger grid spacing of 5km to avoid sampling at a higher resolution than the originals. Grid extents were chosen such that the new grid lay entirely within the original one, as illustrated in Figure 5.7. The desired longitude/latitude grid points were generated and transformed to rotated pole coordinates, resulting in a set of points that no longer lay on a rectangular grid. The regular grid input data in rotated coordinates was interpolated at these points using the `RectBivariateSpline` from SciPy, (optimised for when the input data is on a regular grid). The resulting data point matrix was finally associated with the new regular grid in standard coordinates and archived as NetCDF4 files, one file per variable per year. This was performed using the UCL Legion high performance computing (HPC) service.

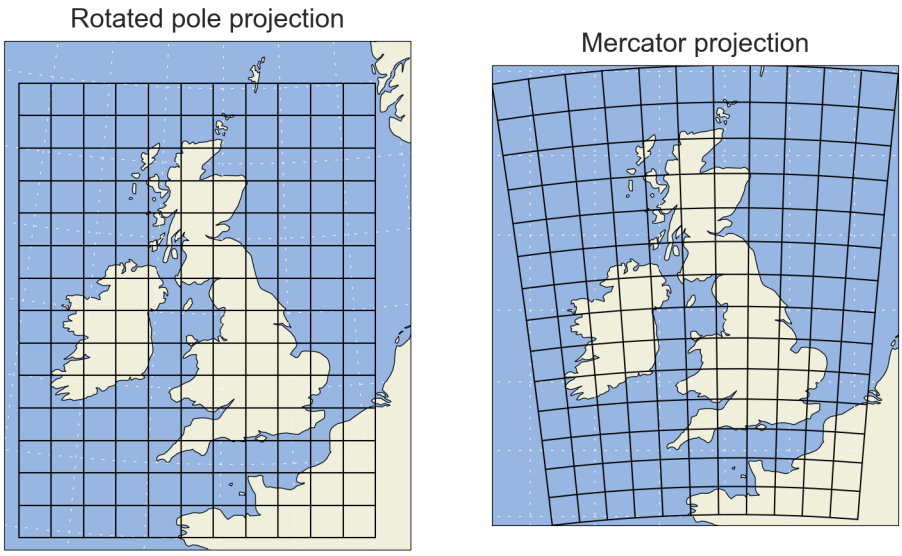

Figure 5.6: Left, an illustration of the original MetOffice grid (with reduced grid density for clarity) on the rotated pole coordinate projection. Right, the same grid shown on a Mercator projection in longitude/latitude coordinates. Note how the perpendicular grid lines in the original projection become curves when reprojected onto the standard coordinate system

5.5 Dwelling monitoring data

5.5.1 EHS and EFUS data

The English Housing Survey (EHS) is a continuous national survey commissioned by the Department for Communities and Local Government (DCLG). It collects information about people's housing circumstances and the condition and energy efficiency of housing in England (DCLG 2016). This research used data from the 2011 Energy Follow-Up Survey (EFUS), which is a sub-sample of 2,616 households from the 2010-2011 EHS and includes additional on-site interviews and inspections. For a further sub-sample of 823 dwellings, temperatures were monitored at 20-

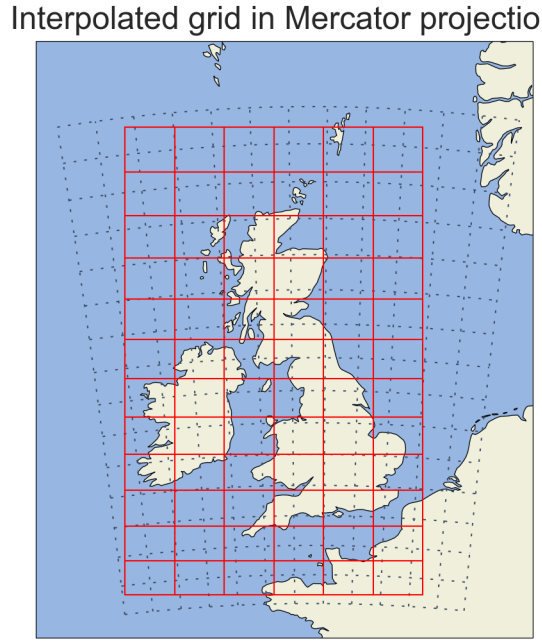

Figure 5.7: Illustration on a Mercator projection map of the interpolated regular grid (red) over the MetOffice rotated pole grid. The new grid has regular spacing and perpendicular lines in the standard longitude/latitude coordinate system

minute intervals in three zones within the dwelling, (living room, hallway and main bedrooms where available) for approximately one year (DECC 2013b). These datasets are of interest as they are designed to be nationally representative for England. However, energy data consists only of electricity yearly readings, precluding its used for thermal modelling. Nevertheless, the internal temperature data collected is useful for developing the modelling of the relation between indoor and outdoor temperatures as described in Chapters 3 and 4, and the associated dwelling characteristic information (dwelling age, type etc).

For the EFUS sites, temperature loggers recorded temperatures at 20-minute intervals over a period from the beginning of February 2011 to the end of January 2012, capturing a full winter and summer period. Outdoor temperatures were not collected using the temperature loggers during the temperature survey, instead hourly data from the nearest Met Office weather station system (Met Office Integrated Data Archive System or MIDAS) was used (MetOffice 2013). It was reported that 75% of dwellings were within 24 km of a weather station, with the mean being 11 km (Hulme, Beaumont, and Summers 2013b). Temperatures were resampled to hourly averages to match the weather data. This was filtered by removing points where indoor temperatures were below 0 $^\circ$ C or above 40 $^\circ$ C, under the assumption these were erroneous - the number of excluded points is shown in Table 5.3.

Table 5.3: Number of hourly temperature readings from EFUS excluded by the data temperature cut-off filters.

Sensor location	N > 40 °C	N < 0 °C
Bedroom	24	15200
Hallway	80	116534
Living room	63	71559

The hourly temperatures were then converted to daily averages and the dwelling mean temperature was calculated as the mean of the bedroom, hallway, and living room. In the case that data was missing from one or more columns, the mean of the remaining columns was used.

5.5.2 Energy meter datasets

Four smart meter datasets were obtained - the Energy Demand Research Project, EDF subset of EDRP (EDRP-EDF), Solid Wall Insulation (SWI) field trial, and the Pennyland Project, each offering a different mix of benefits and drawbacks. It should be noted that the meters used did not meet (and in fact somewhat pre-date) the official smart meter specification discussed in Section 2.4. Table 5.4 offers an overview of the dwelling datasets used. Broadly, the larger the dataset the less detailed information was available, whether this concerned location or site metadata. Consequently, it was important to make use of this range of datasets to address different aspects of this research. The EDRP-EDF and SWI datasets were used particularly in the method development and evaluation phase (Chapter 6) while the EDRP, EDRP-EDF and Pennyland Project were used in the Results Chapter 7.

Table 5.4: Comparison of energy metering datasets. See also Figure 5.12 and Figure 5.13 for plots of dwelling type and age distributions, where relevant. †For sites retained after data cleaning

Variable	EDRP	EDRP-EDF	SWI	Pennyland
No. sites	18370	1879	86	76
No. sites after cleaning	14598	1872	77	76
No. sites with data	14598	1872	69	76
Start date	2008-01-08	2008-01-08	2010-01-18	1980-10-08
End date	2010-09-30	2010-10-26	2012-03-01	1983-03-09
Metered frequency	1/2 hour	1/2 hour	15 minute	Week
No. days of data	8701447	1128694	52328	-
No. weeks of data	1243063	161242	7475	9652

Variable	EDRP	EDRP-EDF	SWI	Pennyland
Has dual fuel data	8466	1048	53	76
Has dual fuel data† (%)	58	56	77	100
Internal temperature data	No	No	Yes	Yes
Location identifier	NUTS1	Postcode district	Postcode	Postcode
Location precision	Low	Medium	High	High
Geographic distribution	UK	Southeast England	England	Milton Keynes
RdSAP survey	No	No	Yes	No
Dwelling age	No	Yes	Yes	Yes
Dwelling type	No	Yes	Yes	Yes
No. rooms	No	Yes	Yes	No

The Energy Demand Research Project (EDRP) was a major project from 2007 to 2010 implemented by 4 major energy suppliers to test smart metering infrastructure and measure customer response to energy feedback. Each supplier conducted similar (but not coordinated) trials, precluding comparisons between trial groups. An extensive overview of the trial design, execution, and results can be found in the final analysis report of the project (Raw, Ross, and Cripps 2011). The 18,370 dwellings involved were distributed across England, Scotland, and Wales (Figure 5.8). The large number of participants made this an interesting dataset, which additionally was representative of the data expected to be collected from the ongoing rollout of smart meters, since meters matching the official technical specification were used. The limited contextual and metadata provided with the public release EDRP dataset was its biggest limitation - the anonymisation process removed nearly all site specific information. Site location data was provided at a relatively coarse level, which limited the accuracy of the weather data lookup by requiring weather to be averaged over a larger area (see Section 5.6).

The UCL Energy Institute obtained EDRP data from EDF for 1,879 properties, with additional metadata relative to the public release EDRP dataset. These sites were located only in London and Southeast England (Figure 5.9). Unfortunately, these sites cannot be cross-referenced with the rest of the EDRP dataset because the customer IDs have not been preserved across the datasets. The additional metadata fields include basic information about the dwellings that could help understand their performance - notably the number of rooms and occupants, the dwelling age and type, and the postcode district which enables precise geolocation.

The Solid Wall Insulation (SWI) field trial was a project conducted by the Energy Saving Trust (EST) from 2010 to 2012 to investigate the in-situ performance of various forms of solid wall insulation. Solid-walled dwellings (including stone and non-cavity brick walls) are both particularly inefficient and difficult to insulate and therefore warrant special attention. The detailed site monitoring included:

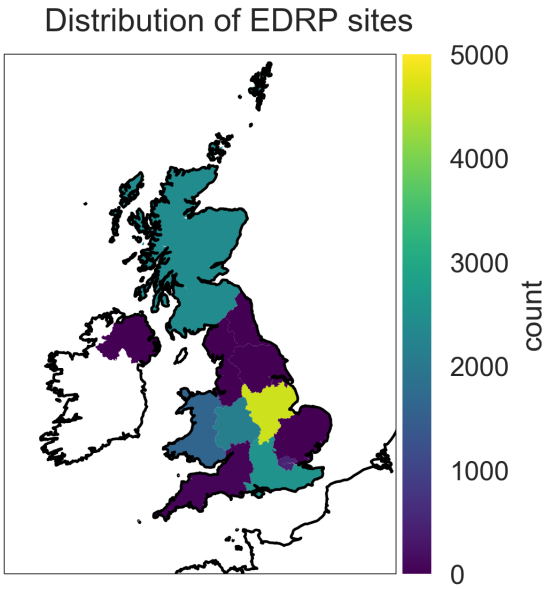

Figure 5.8: Geographical distribution of EDRP sites across UK NUTS1 regions

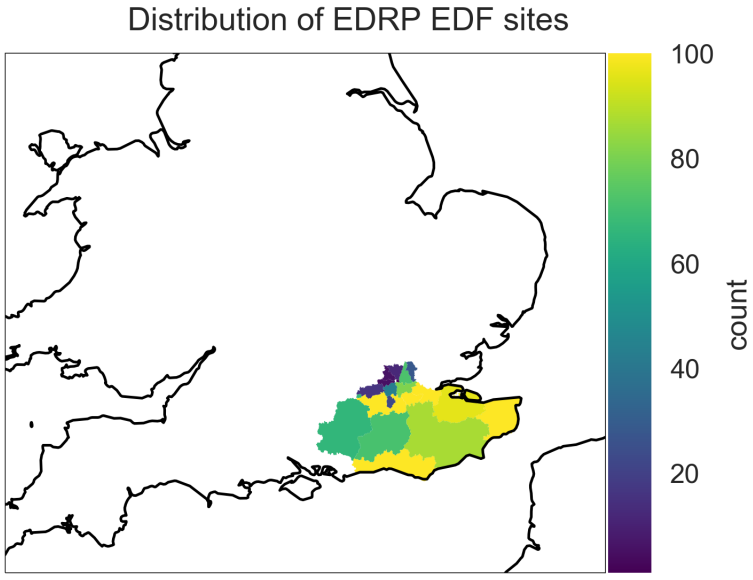

Figure 5.9: Geographical distribution of EDRP-EDF sites across postcode areas in South-East England.

- Energy monitoring using smart metering devices on the electricity and gas supplies.
- RdSAP surveys with recording of inputs for the RdSAP2005 calculation.
- Architectural surveys and plans (more in depth than the RdSAP survey).
- Thermal imaging (with the aim of performing a before/after comparison).
- Air-tightness testing for some sites.
- Internal temperature and relative humidity monitoring in several rooms.
- For some sites, heat flux sensors measure the U-values of the outside walls.

The internal temperature data was of particular interest, enabling notably in the context of calculating C_T, this is explored in Section 6.2. SWI only contains a small number of dwellings with a specific construction that only represents 30% of UK dwellings. This dataset also had major quality issues in terms of organisation, errors, and missing data.

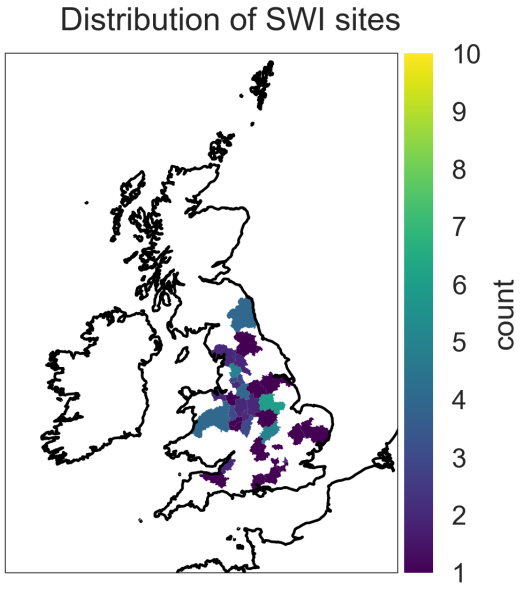

Figure 5.10: Geographical distribution of SWI sites across postcode areas across England and Wales.

The Pennyland Project was a large-scale housing project carried out in Milton Keynes by members of the Open University Energy Research Group between 1976 and 1984. The overall aim of the project was to investigate the energy savings achievable through better house insulation and the incorporation of passive solar design features in new buildings. A housing estate was constructed with a small number of dwelling designs intended to test insulation and passive solar features compared to more classic designs, in order to evaluate the resulting differences in energy consumption (Chapman, Lowe, and Everett 1985). Of 177 houses in the original study, data for 76 remains accessible today. Energy and temperature readings were taken on a weekly basis using conventional monitoring equipment (not smart meters).

While not a smart meter dataset, Pennyland is relevant to this research for a number of reasons. Although not providing daily averages, the availability of weekly data can enable Deconstruct to

be applied. The fact that the dataset includes a known set of different building types of very specific construction offers a rare opportunity for evaluating Deconstruct by determining whether it is able to distinguish between the different building types. Importantly, Pennyland properties were extensively studied during their design, construction, and operation.

In both the EDRP and the EDF datasets there is a gap between the number of customers that should have both electricity and gas data and the number for which at least one record of each is found. The EDF dataset in particular shows that while most of the customers in the test set were 'dual fuel' customers (getting both electricity and gas from EDF) a significant proportion of those appeared to have no gas data. It was possible for sites to use supplementary energy sources (such as fireplaces, solid fuel stoves, paraffin burners), but only the SWI dataset contained some indication of which sites made use of supplementary heating.

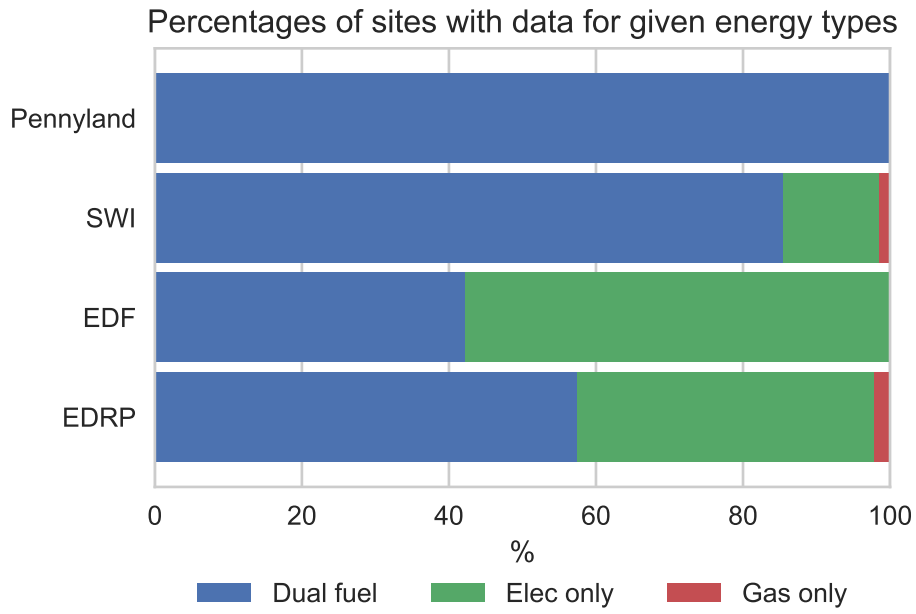

Figure 5.11: Share of energy data types available for sites for datasets with appropriate metadata

The EDF and SWI datasets had dwelling type information, while Pennyland sites were mostly semi-detached dwellings built at a similar time. Only a small fraction of the sites in EDF or SWI datasets were classed as flats compared to the national average, the rest were various forms of detached or semi-detached houses, with the SWI set containing a larger fraction of terraced houses compared to both the EDF set and the national average, see Figure 5.12. The differences between EDF and SWI datasets are much clearer when looking at age bands Figure 5.13 - SWI sites are exclusively solid-wall construction pre-1945 buildings which pre-date modern construction methods and insulation, one site dated from 1496. The EDF sites meanwhile are in line with national averages. House age and type information was not available to for the full EDRP dataset.

The daily power demand distributions of the EDRP, EDRP-EDF and SWI sites across months, split by weekday/weekend are shown in Figures 5.15, 5.16, 5.17 (filtered for outliers in all cases).

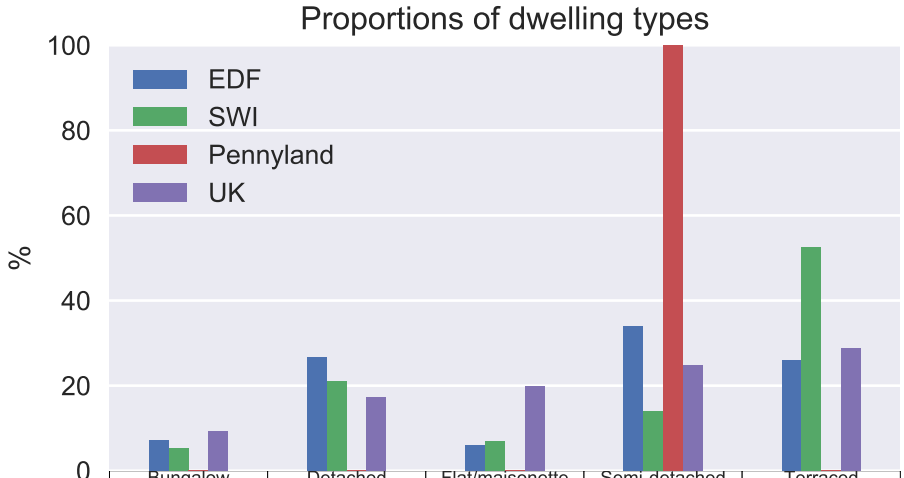

	Bungalow	Detached	Flat/maisonette	Semi-detached	Terraced
EDF	109	398	91	510	390
SWI	3	12	4	8	30
Pennyland	0	0	0	76	0
UK	2147000	4044000	4660000	5791000	6730000

Figure 5.12: Percentages of house types in the EDF, SWI, and Pennyland datasets compared to the UK distribution for 2014 DCLG (2013)

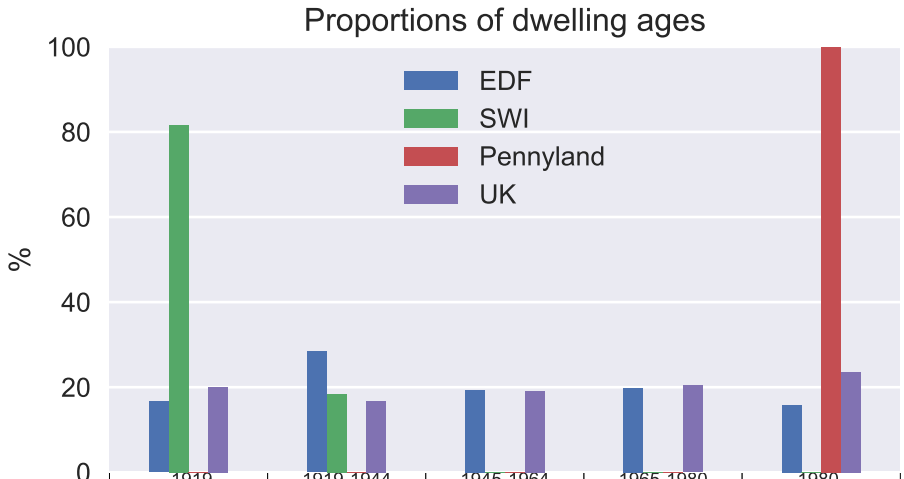

	-1919	1919-1944	1945-1964	1965-1980	1980-
EDF	242	413	282	288	229
SWI	62	14	0	0	0
Pennyland	0	0	0	0	76
UK	4698072	3902999	4473141	4774762	5521977

Figure 5.13: Percentages of sites in age bands for EDF, SWI, and Pennyland datasets compared to the UK distribution for 2014 DCLG (2013)

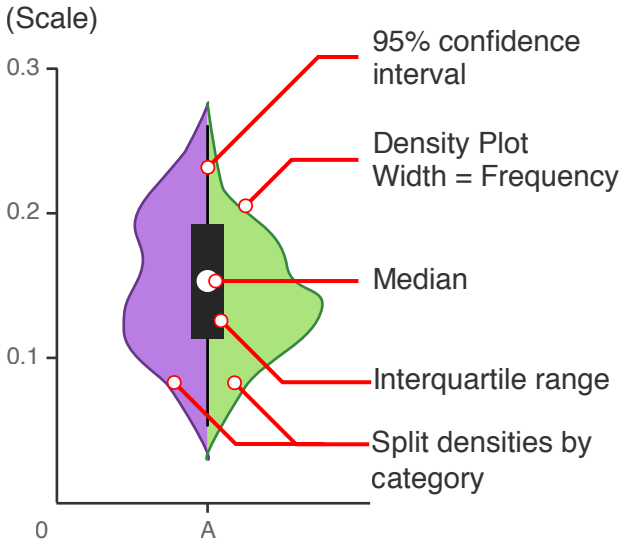

Figure 5.14: Explanation of Violin plot. Densities are estimated using a Gaussian kernel density estimator, with plot densities optionally split left/right by a categorical variable (such as weekday/weekend for daily energy data) ("Violin Plot" 2017)

A Violin plot is chosen to visualise the distribution of the data, as well as the median and interquartile range, as it can summarise many aspects of the power distribution within one plot. For a description of the violin plot elements, see Figure 5.14.

It is important to note that the smaller EDF and SWI datasets display more weekday/weekend asymmetry in summer than the full EDRP dataset, suggesting some smoothing occurs from the larger aggregation. The SWI dwellings consume considerably more energy overall than the EDRP dwellings, even in summer.

5.5.3 Data acquisition and extraction

This section presents the data acquisition and extraction of the datasets used. Each dataset was received in a different format, and therefore presented different processing challenges. Each dataset was extracted using a bespoke Extraction, Transform, and Load (ETL) workflow. These are presented in this section. Energy time series data presented particular challenges. Identification and tracking missing values and errors was poorly documented, and when performed often did not account for all errors eventually found. There was a similar lack of consistency in formatting of time stamps and accounting for time zones and daylight savings time changes resulting in forward/backward jumps when clocks changed. This thesis uses UTC (Coordinated Universal Time) with time zone information, which all datasets were converted to.

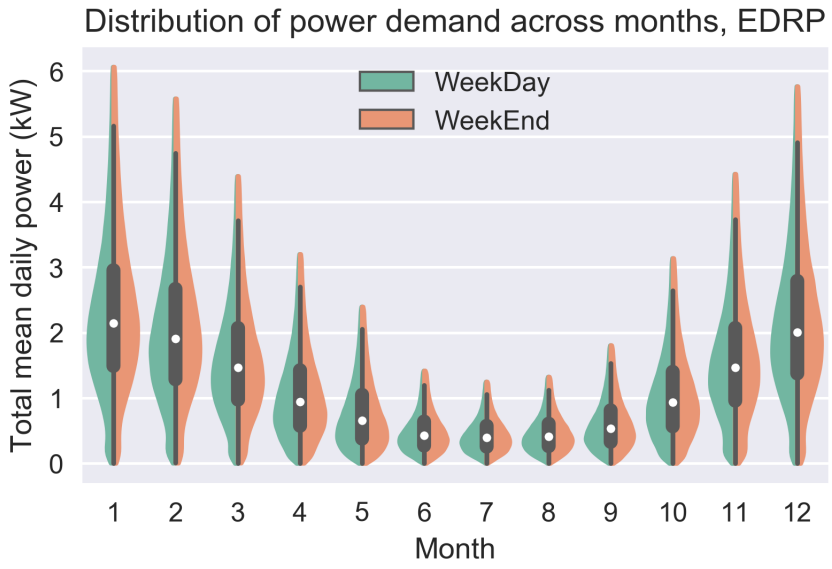

Figure 5.15: Violin plot (showing kernel density estimate of distribution) of daily mean total power demand (electricity and gas) for EDRP sites across months, with differences between weekday and weekend.

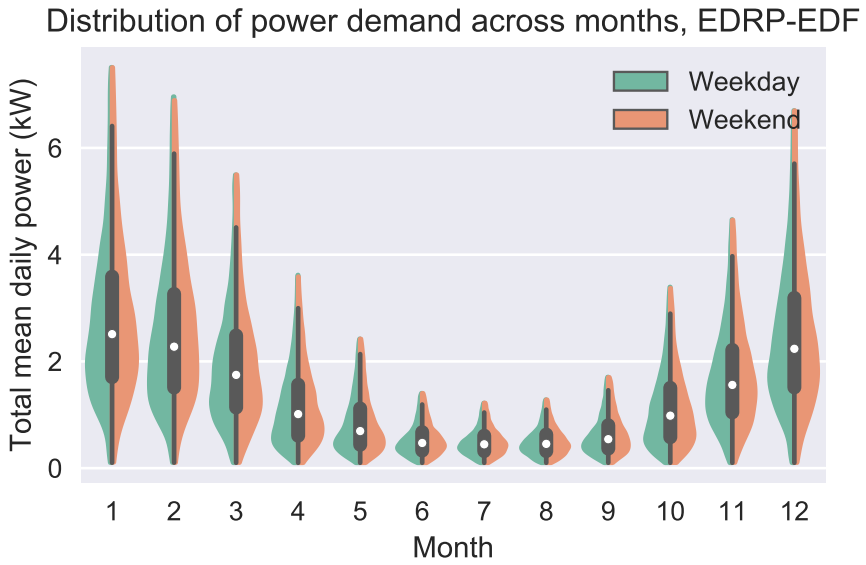

Figure 5.16: Violin plot (showing kernel density estimate of distribution) of daily mean total power demand (electricity and gas) for EDRP EDF sites across months, with differences between weekday and weekend.

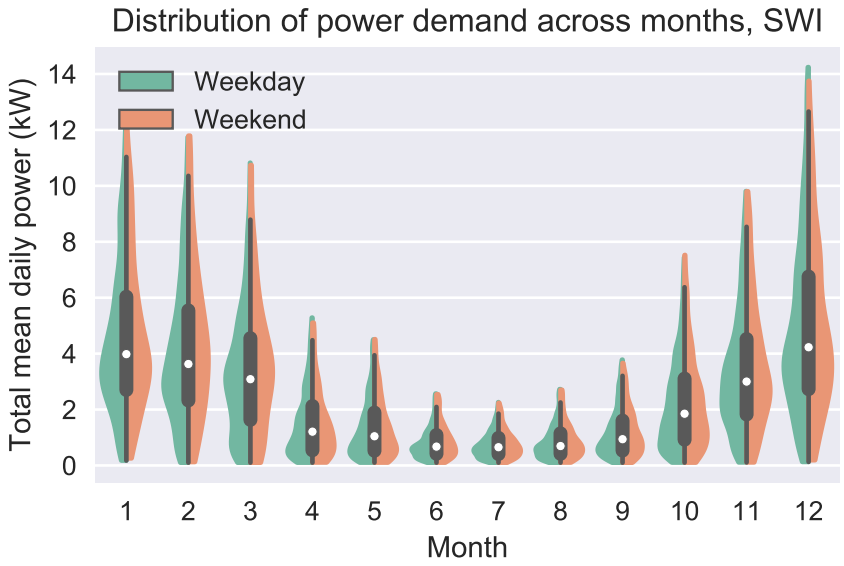

Figure 5.17: Violin plot (showing kernel density estimate distribution) of daily mean total power demand (electricity and gas) for SWI sites across months, with differences between weekday and weekend.

5.5.3.1 Energy Demand Research Project

The EDRP public release data was obtained from the UK Data Service repositories (UK Data Service 2014a). This data was consolidated from the four energy suppliers' data and anonymised and de-duplicated following rules described in UK Data Service (2014b). The number of sites from each supplier is shown in Table 5.5. Input files consisted of Comma-separated values (CSV) tables for electricity meter readings, gas meter readings, geographical site information and other site metadata. The electricity file contains 413,836,038 readings and was 12GB in size, the gas file contains 246,482,700 readings and is 9GB in size. The CSV files were imported into the SQL database following the workflow shown in Figure 5.18 which included the following steps:

- The electricity and gas half hour meter readings were imported to a temporary table
- The original text time stamps were parsed into a true date-time format and time zone information was added to account for daylight saving time.
- The site metadata table was created by joining the metadata tables for the electricity, gas, and geographical information.

Table 5.5: Breakdown of numbers of participants in EDRP study by participating energy utility company (Raw, Ross, and Cripps 2011)

Supplier	Smart meter customers
EDF	1879
E.ON	8055

Supplier	Smart meter customers
Scottish Power	1330
SSE	7106
Total	18370

5.5.3.2 EDF subset of EDRP

EDF's EDRP data was drawn from EDF customers sampled using the strategy described in Raw, Ross, and Cripps (2011). It aimed to minimise confounding factors, e.g. by excluding EDF staff and customers on special tariffs. The sample was stratified according to the estimated annual energy consumption (low, medium, high bands).

One CSV file per experiment group per energy type was provided. Most sites show systematic data gaps around the 10th August 2010, which was assumed to be a problem with the original data collection. The files were processed according to the workflow shown in Figure 5.19 to produce combined half-hourly time series tables for electricity and gas, as well as a site metadata table. For each file:

1. Columns were renamed to a consistent format
2. Dates were converted from raw text to timestamp objects
3. Duplicate rows were dropped
4. Sites which still contained non-unique (site, reading date) pairs after row de-duplication were dropped.
5. The 48 columns corresponding to half-hourly readings for the day row were pivoted to a 'long' format with one row per reading and a new date-time column was added with time zone support for daylight saving time
6. The daily average table was created from the half hourly table
7. Site metadata was extracted to a separate table. Where metadata existed from both electricity and gas tables these values were merged ensuring that fields matched where expected.
8. The provided postcode fragment (first 4 characters with no spaces) was processed to produce valid district codes. Certain postcode fragments were ambiguous - for example the fragment NW23 could correspond to NW1 6XX or NW16 4XX (district code NW1 or NW16). These were first filtered using the postcode format rules discussed in Section 5.3.1. The remaining ambiguous cases which had a valid code format but did not correspond to real places were cross-referenced with the known-real code list.

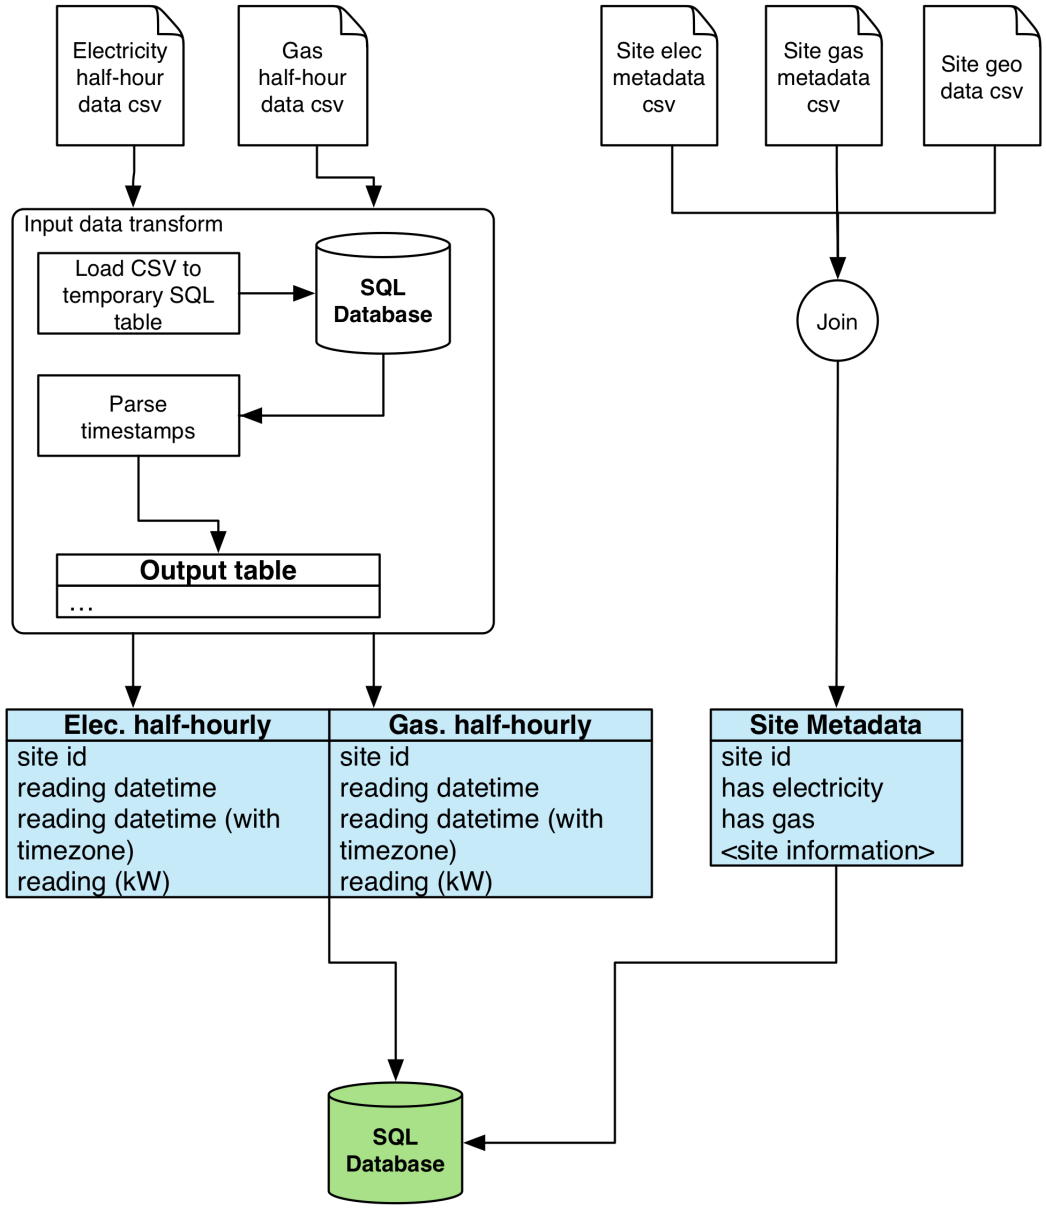

Figure 5.18: Summary of EDRP extraction and cleaning process.

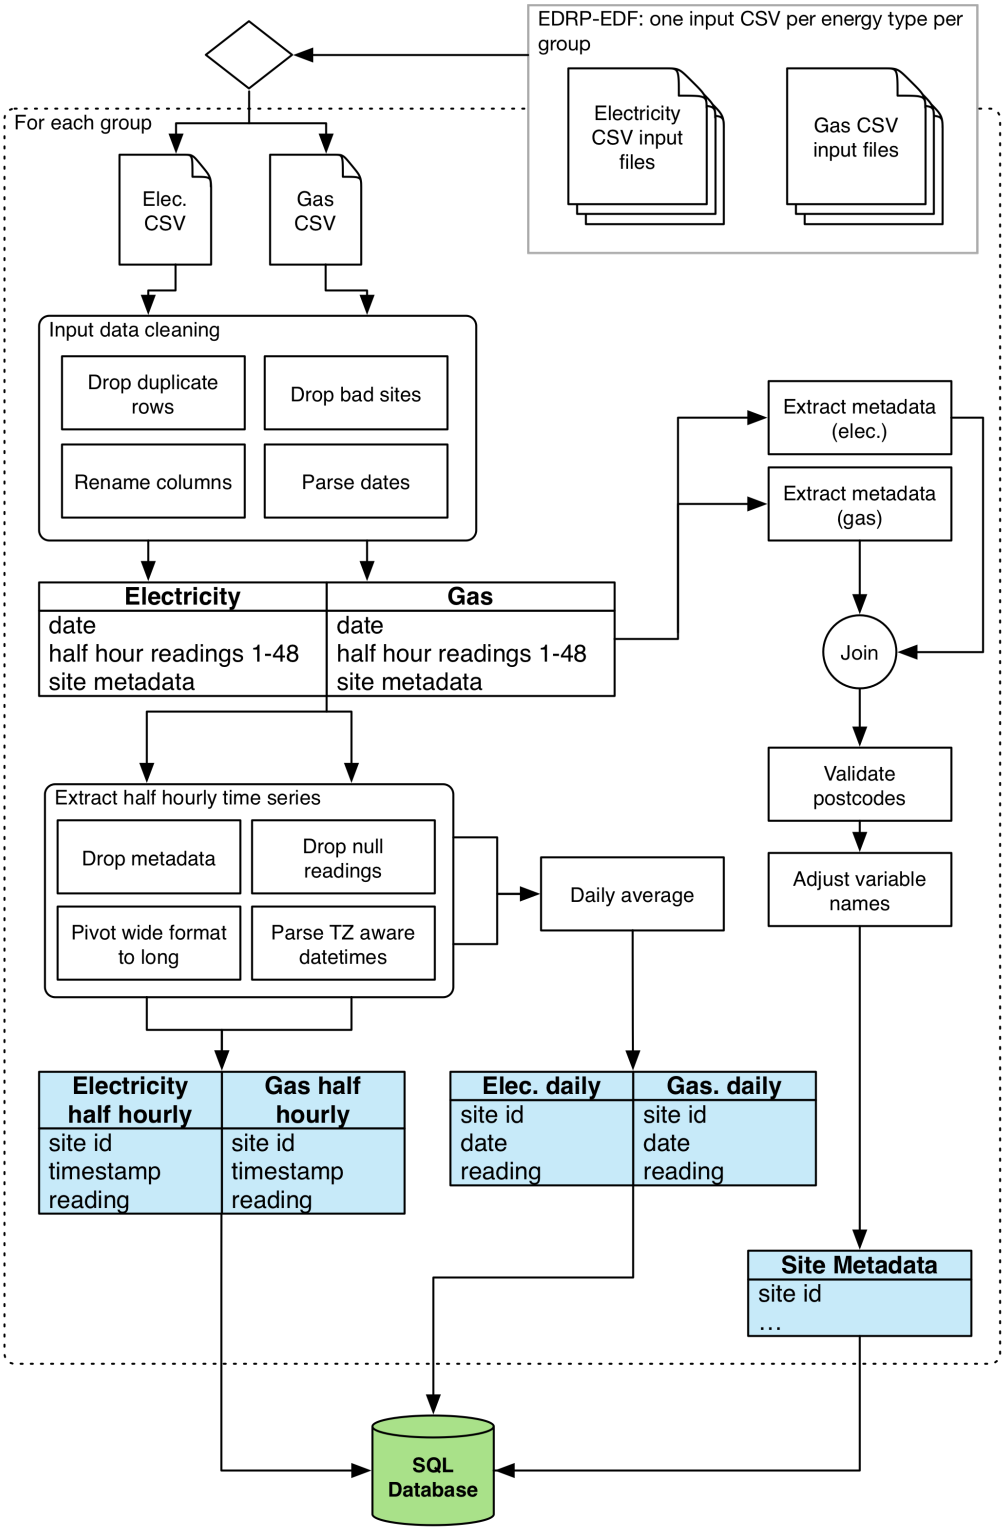

Figure 5.19: Summary of EDF subset of EDRP extraction and cleaning process

5.5.3.3 Solid Wall Insulation field trials

Extensive descriptions of the project and sites can be found in the project report (Birchall, Pearson, and Brown 2011). The monitoring and record keeping for these trials was generally inconsistent, with evidence of a number of hardware failures as well as a lack of consistency in the data storage. Data delivered included:

- SAP reports, site surveys, and air tightness tests as Portable Document Format (PDF) files.
- Thermal imaging and site photographs.
- Energy and temperature as a mix of CSV and Excel files with various layouts. Energy values were recorded as cumulative meter readings in kWh for electricity or m^3 or ft^3 for gas (unlike other datasets where gas data was delivered in terms of gas energy content).

The diversity of file types, layout within folders, and file formats posed a particular challenge to a reliable automated data extraction. An overview of the process can be seen in Figure 5.20.

Table 5.6 summarises sites omitted during import. Site metadata was extracted from pdf files using several pdf scraping tools. Readings for site 41 between 2011-02-04 12:30 and 2011-10-01 00:00 were missing the decimal point in the CSV file, this was repaired manually. Site location in longitude/latitude coordinates was determined by geocoding the postcodes using an online service (Bell 2016).

Table 5.6: Summary of sites omitted

Site id	Comment
17	no SAP data
19	oil fuel heating
56	two houses (56a, 56b) merged
106	oil fuel heating
112	solid fuel heating

Unlike the other datasets, gas consumption in SWI was delivered as volume readings instead of energy. Gas energy consumption in kWh was derived from gas volumetric consumption meter readings. First, the mix of imperial and metric readings (ft^3 and m^3) were converted to metric (including accounting for sites which switched part way through). Many gas meters were reset or replaced during the course of the experiment. It was assumed that the meter had been reset if the reading at a given time stamp was less than 1% of the reading at the previous time stamp. Using a percentage change for detection rather than searching for a reset to zero was necessary for two reasons: firstly, some gas may be consumed between the reset time and the next reading and secondly, some of the meters showed 'jitter' in their readings meaning they could decrease slightly

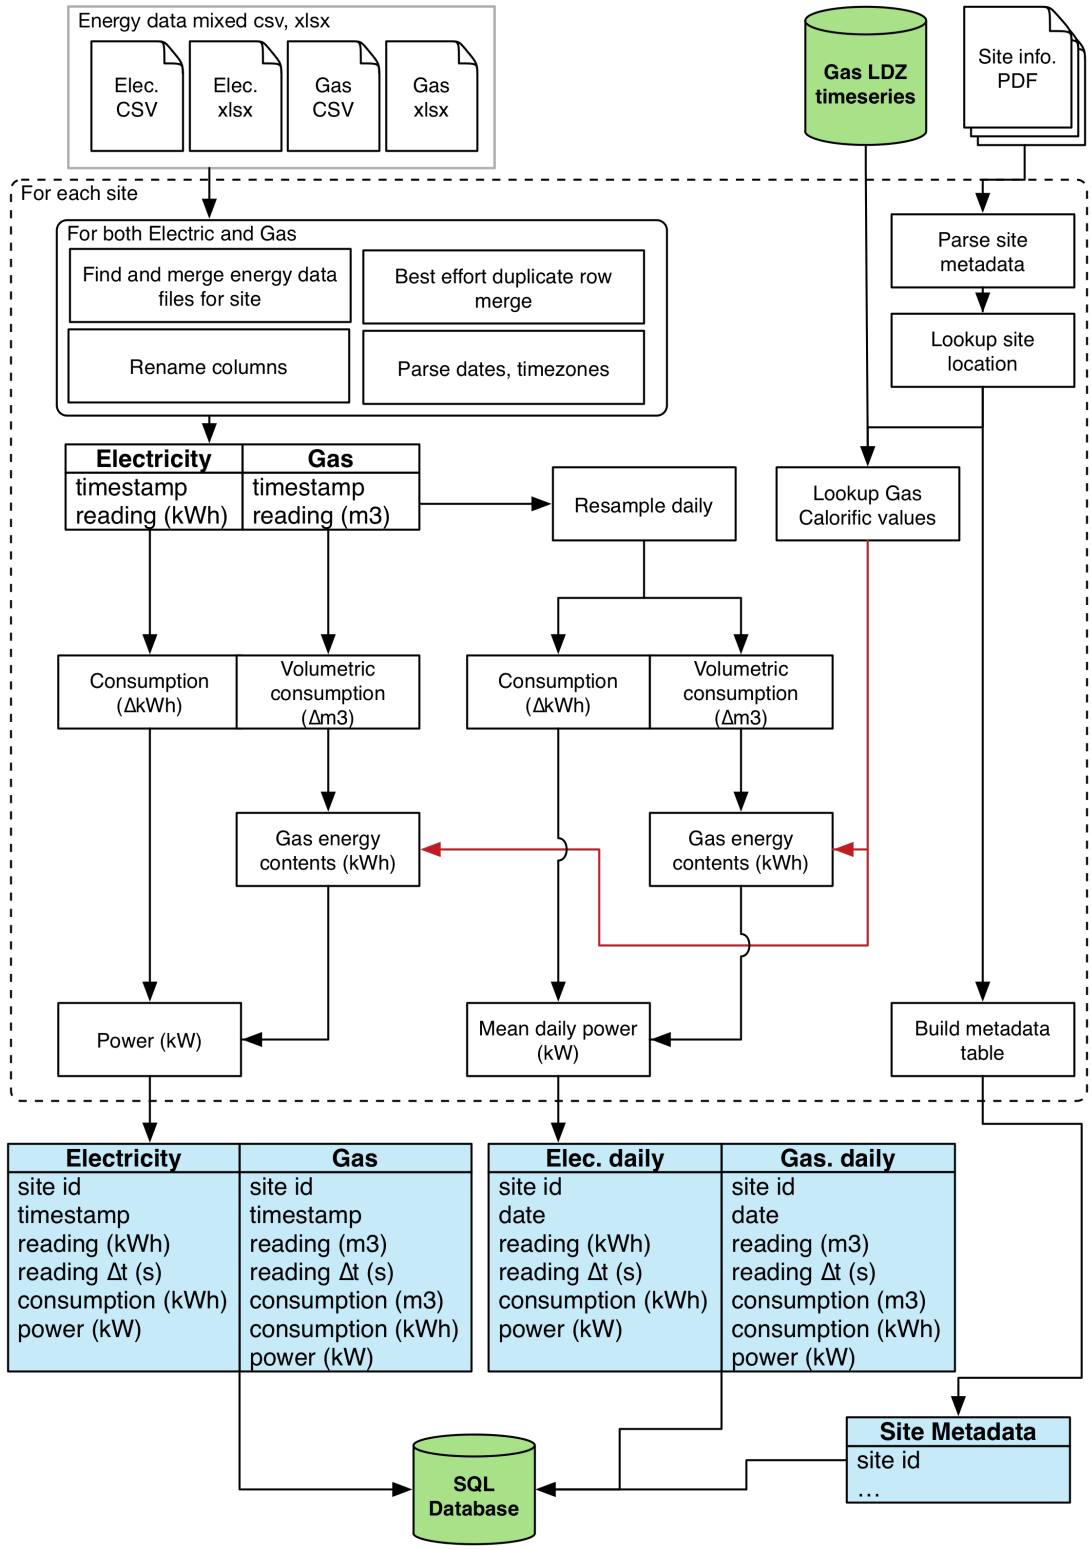

Figure 5.20: Summary of SWI trials data extraction and cleaning process

between readings.

The National Grid reports daily calorific values for each Lower Distribution Zone (NationalGrid 2015), which was downloaded for the days covering the SWI monitoring period. Daily average CV time series were obtained for the gas Lower Distribution Zone (LDZ) corresponding to the site locations. Gas volumetric consumption for each time step was converted to energy content using the CV for the corresponding day and LDZ using eq. 3.15. As discussed in Section 5.3.2, the difference in CV between zones was relatively small (usually less than 3%) but still sufficient to motivate using the day and location specific CV.

The electricity and gas cumulative energy use in kWh was converted to power in kW. Energy data files for each site were imported and merged into single time series tables with suitably parsed time-stamps. Power (P) was found by calculating the energy consumption E per time step dt and the size of each time step (which was not guaranteed to be regular) and setting P=E/dt for every time step with a valid consumption value.

Daily average power values were obtained using the energy/time delta between total energy readings (kWh) at midnight on each day. This was preferred to averaging the original powers as it avoided accumulating measurement errors from the meters and reduced the percentage error in the reading delta. The resulting time series had a large number of gaps (before any additional outlier filtering) due to issues in original data collection (Figure 5.21).

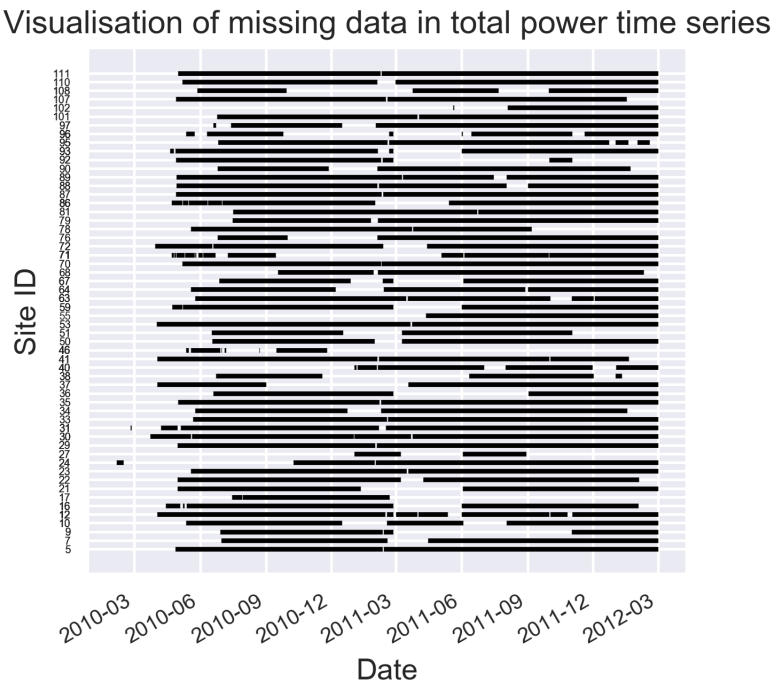

Figure 5.21: Visualisation of the missing data in the SWI sites total power time series. Dark lines represent the time series for each site, gaps in the lines represent missing data for that site.

The SWI trials included monitoring of domestic internal temperatures which displayed similar

issues in data and metadata collection to the energy data. Sensor ID codes were not associated with a room within the dwelling, but it was known that two or more internal measurements and one external measurement were monitored. The external sensor was assumed to be the one with the lowest temperature readings in the winter months, while mean internal temperature was calculated using the mean of the remaining sensors. A range of issues in the data were manually corrected, notably, site 111 was excluded because the monitoring files were completely inconsistent, while for site 33 one sensor gave readings in excess of 500 $^\circ$C, this sensor data was discarded and the remaining sensors used.

5.5.3.4 The Pennyland Project

Weekly energy from the Pennyland project was retrieved from the UK Data Service archives. These files were transcriptions of the original tape format archive. The archival data had been extensively pre-processed with errors and missing values identified and the data laid out in a consistent table. Recovering the data required the development of a simple parser for the legacy data format based on PDF scans of the original manuals. The weekly readings were uploaded to the PostgreSQL database with a similar table format to the smart meter datasets. No site metadata was included within the data tables. Instead, project descriptions from Chapman, Lowe, and Everett (1985) were used. Weather data was associated with the sites by taking the latitude/longitude location of the energy park and retrieving historical climate data from the CFSR project, which had a much longer time coverage than the MetOffice data.

5.6 Data loading and linking

Site data, energy data, and weather data were imported into their respective data stores in a homogenised form, enabling a common process to be used to load linked energy and weather time series for analysis. Figure 5.22 shows an overview of this process. Starting with a site metadata entry, energy data from that site is loaded from the PostgreSQL database. In parallel, location data from the site metadata is used to run a PostGIS query which determines which weather grid polygons correspond to the site locations and loads the corresponding weather time series.

For some sites only regional rather than point location data was available, therefore several weather grid units may be returned since the region polygon (e.g. a postcode area) may overlap with several weather grid polygons. In this case, the average of the time series for those squares is taken, weighted by the fractional area by which the grid square overlaps the location polygon. An illustration of the process is shown in Figure 5.23.

The weather grids in netCDF files are indexed by longitude/latitude points - that a time series can be fetched given the longitude/latitude point on the grid. To facilitate geographical queries,

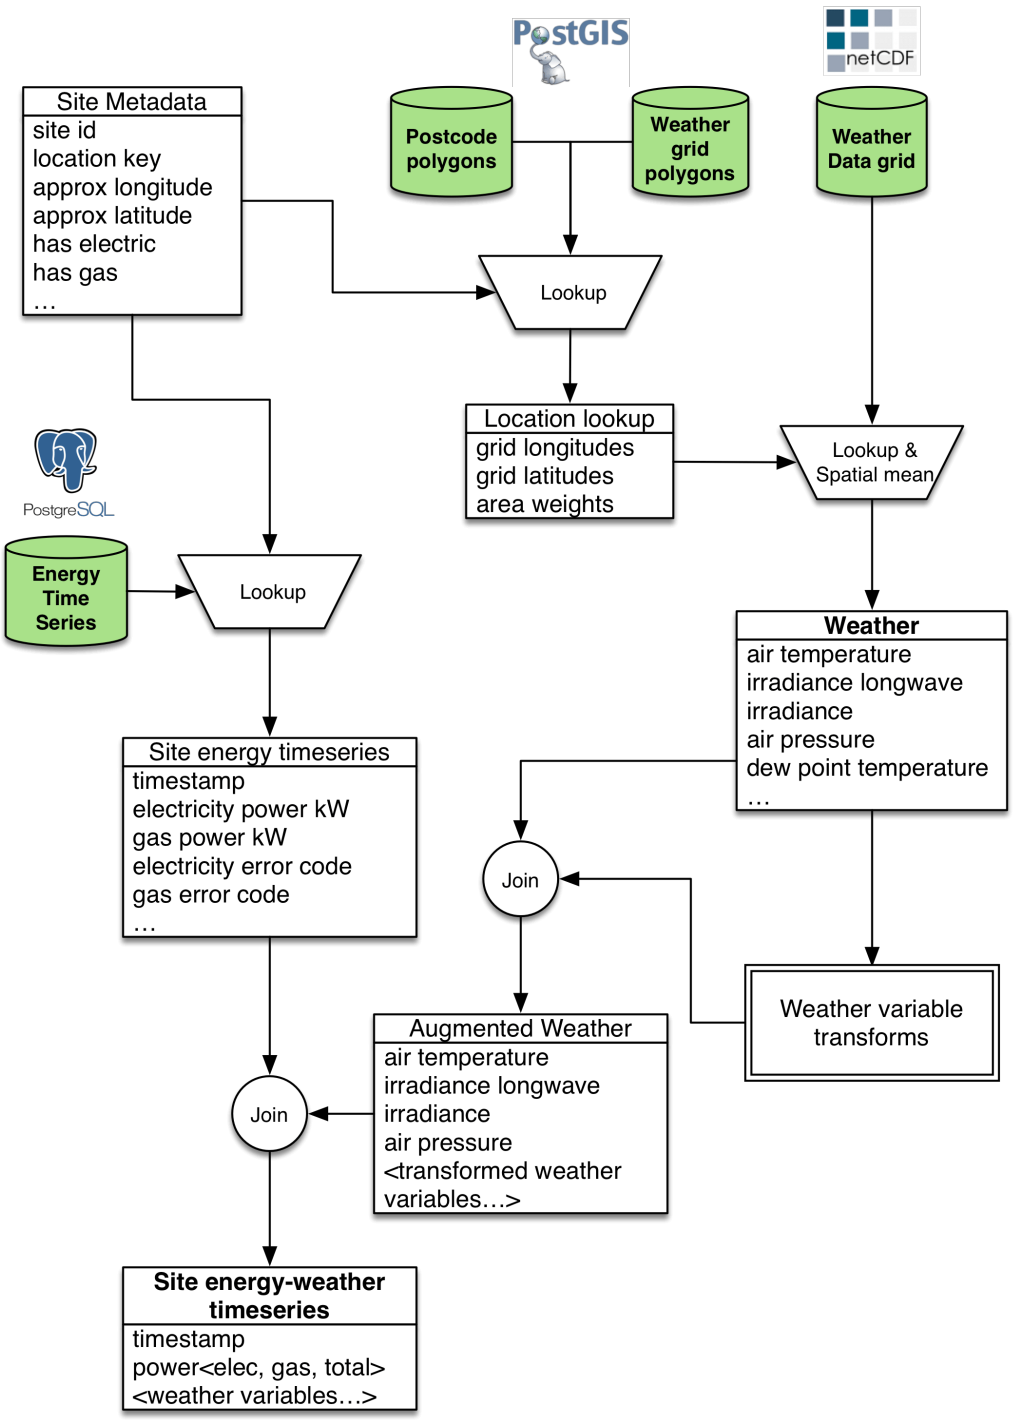

Figure 5.22: Top-level overview of process for loading energy/weather time series datasets from using the site data, smart meter data, and weather data previously imported to the data stores

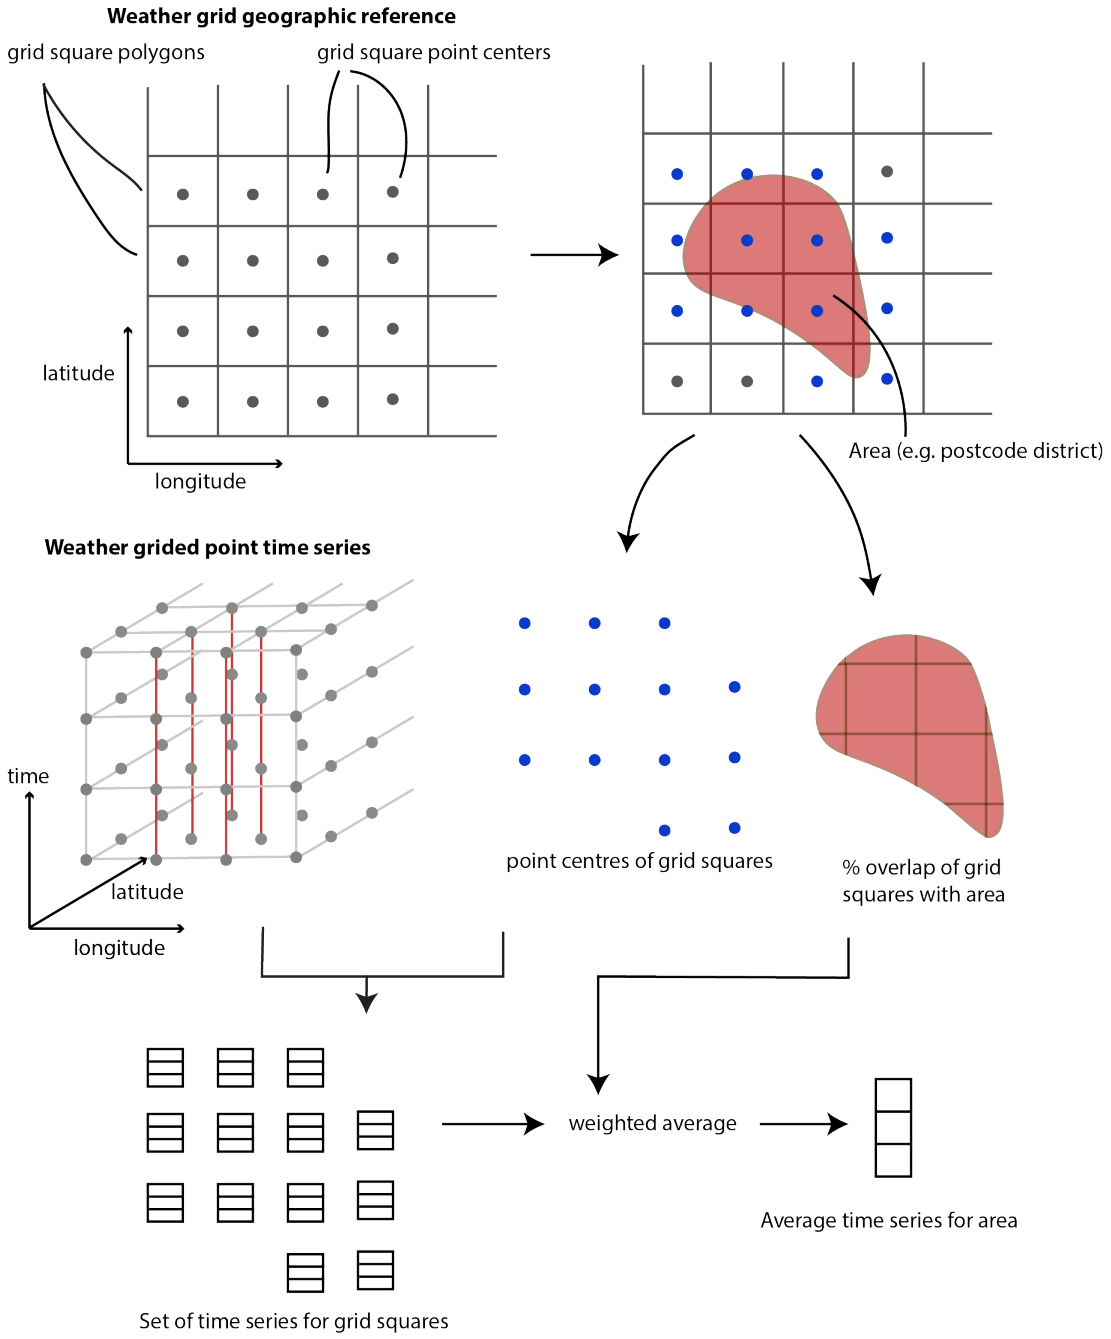

Figure 5.23: Illustration of spatial averaging of one weather variable time series using a lookup of the overlap of a geographic region, such as a postcode district, and a grid of polygons representing the weather data grid to determine the correct indexes for weather time series from netCDF files

a PostGIS table was created for each weather dataset containing points for each weather grid longitude/latitude point and polygons representing the grid area surrounding each point. This in effect assigned the value of each weather grid point to the surrounding grid area. It was then possible to run a query to find the overlap between a given area polygon, such as a postcode district, and the grid squares. This operation returned the overlapping grid squares, the points at the centres of those squares, and the fraction by which they overlapped with the postcode district.

The points were used to lookup weather time series for each variable; as many time series were retrieved as there are grid squares overlapping with the queried area. A weighted average was taken of the time series using the fraction area overlap as weights, which avoided giving full weight to the weather in a grid square that may only overlap very little with the query area.

Optionally, additional weather variable transforms can then then applied. For example, if solar irradiance on an inclined plane is desired, this is calculated from the weather and location data using the method described in Section 3.6.1.1. Finally, the weather time series is joined with the energy time series to form a table ready for analysis.

5.7 Summary

Processing the datasets needed to undertake the research for this thesis was a time consuming and error prone process. It was found that by putting in place extensive automation of the data workflow, errors could be minimised and more easily corrected when identified by re-running the workflow. The tools for handling these datasets have varying levels of maturity. While weather datasets enjoy an established and extensive support in terms of data sharing standards and data processing tools, energy data remains a largely ad-hoc domain. In cases such as the SWI dataset, the level of inconsistency in data collection and archival made automation of the process exceptionally complex. It is to be hoped that, as smart meter data becomes more prevalent and data collection is routed through official channels (such as the planned DCC data gateway in the UK (Beard 2015)) such problems will diminish. There remains a need for open source, community supported software packages for handling smart meter data which could become repositories of best practices. These would also help accumulate algorithms for common tasks such as calculating energy and power from total consumption in a fault-tolerant and uncertainty minimising way.

GIS and weather data was obtained and linked to dwelling energy meter time series and dwelling metadata. obtained from EFUS, EDRP, EDF, and SWI. These extracted and cleaned from a range of input formats and archived in a data storage and loading system based on an SQL database and NetCDF gridded data files. Automated processes were developed and put in place with a view to streamlining data capture and input, and facilitate error checking. Software was written which facilitated the generation of energy plus weather time series, and which made it simple to load site metadata and site time series data for analysis.

The energy and location data input for Deconstruct was obtained from four EDRP, EDRP-EDF, SWI, and Pennyland datasets. Using GIS information, this was linked to gridded weather data. The UK Metoffice provided high resolution temperature and irradiance data to match with energy time series from EDRP, EDRP-EDF, and SWI datasets. The CFSR dataset was used to provide weather data for the Pennyland project. The EFUS dataset provided nationally representative internal temperature monitoring information, to be used for estimating parameters of the Deconstruct internal temperature model.

Chapter 6

Method Evaluation

This chapter evaluates the Deconstruct method developed in Chapter 4. In Section 6.1 the assumptions made in the data quality filters and sampling method assumptions presented in Chapter 4 Section 4.4, notably the low solar gain sampling used to estimate dwelling HLC, are evaluated. In Section 6.2, coefficients for the internal temperature model presented in Section 3.4 are estimated using several approaches and the effect of the coefficient estimate variants on the HLC estimates compared. Section 6.3 presents Deconstruct model parameter uncertainty estimated using a Monte Carlo simulation approach. In Section 6.4, an investigation is made into the effect of the base-load gain parameter assumption on the internal temperature estimate. Gas boiler efficiencies are used to compare HLC and HTC values. Finally, Section 6.5 provides a calculation of the error in power demand prediction made using inferred dwelling thermal parameters.

6.1 Evaluation of data sampling

As described in Chapter 4 Section 4.4, the Deconstruct method employs a number of data quality filters and sampling methods for the purpose of cleaning input datasets and selecting subsets of daily data suitable for the inference of thermal model parameters. These filters and samples depend on a number of simplifying assumptions. The following sections aim to evaluate these assumptions using suitable dwelling data. They aim to support the sampling approach which takes site power data as illustrated in Figure 6.1.

6.1.1 Dependence of base-load on external temperature

Section 4.4.2 and Section 4.5.2 defined the base-load power P_B as being independent from external temperature. To test this assumption, samples were taken from the EDRP-EDF dataset according to the base-load sample method described in Section 4.4.2. Using sites with at least 3 days of

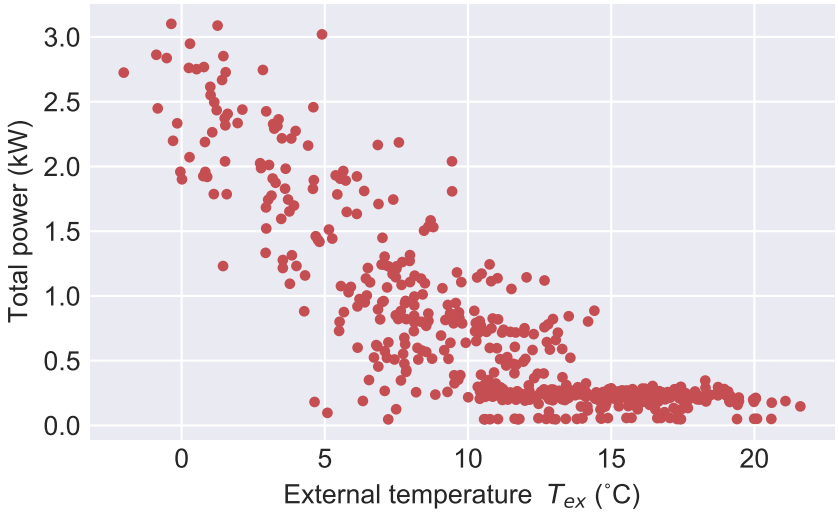

Figure 6.1: Example of unfiltered daily average power-external temperature values from EDRP-EDF site 288

data in the sample, the covariance of power with temperature was calculated using the Pearson correlation coefficient. The resulting coefficients plotted in Figure 6.2 demonstrate that the base-load sample is independent of the external temperature (-0.5 < r-coefficient < 0.5), supporting the assumption that $\frac{dP_B}{dT_{ex}} = 0$ is likely to be valid.

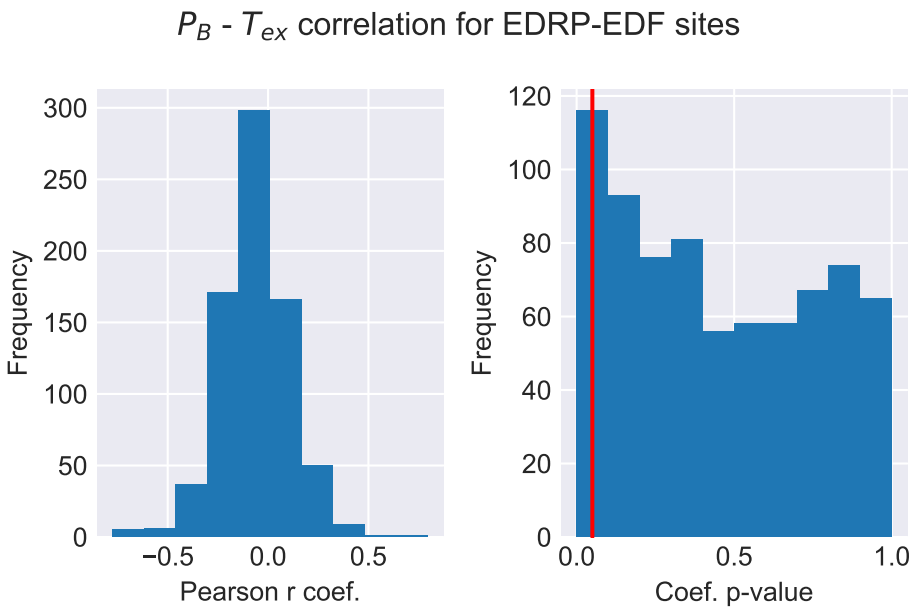

Figure 6.2: Pearson correlation coefficients (left) and coefficient p-value (right) for the base-load sample against the external temperature for EDRP-EDF sites. The red line in the right-hand plot indicates the location of p=0.05

6.1.2 Dependence of power on wind speed

In Section 4.5.2, it was stated that power demand does not empirically display a dependence on wind speed. To test this assumption, the Spearman rank correlation coefficient between total power P_{tot} and wind speed v_{wind} was calculated per site for the EDRP-EDF dataset (Figure 6.3), as well as for a low solar gain sample of site data (Figure 6.4). The Spearman coefficient was preferred over Pearson as wind speeds are not normally distributed. P-value distributions are presented for completeness, although the interpretation of these values in the case of the Spearman test is of only limited utility with sample sizes which are under 500 points. The distributions presented show that the values are not significantly correlated with absolute values under 0.5 (median -0.2 in both cases) (Zwillinger and Kokoska 2000).

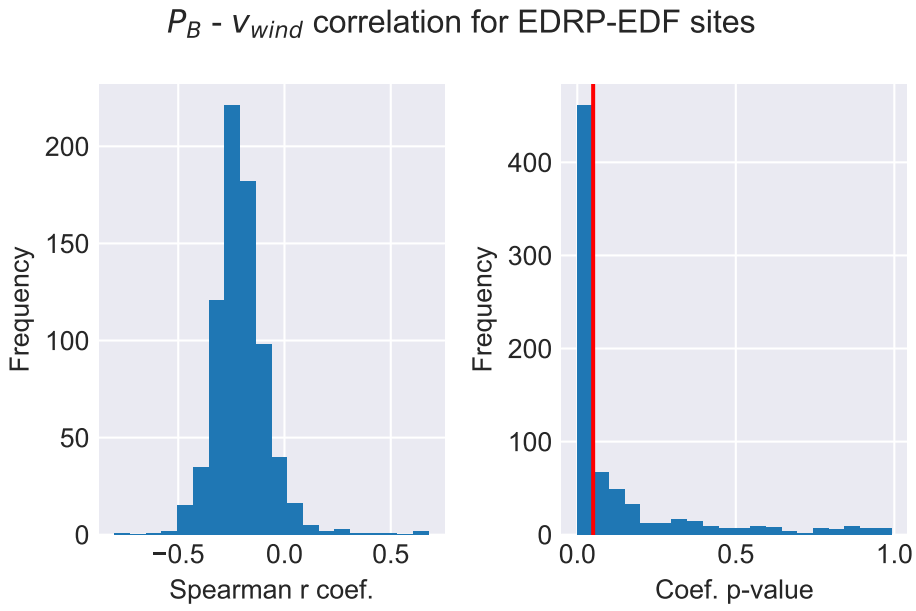

Figure 6.3: Distribution of Spearman rank-correlation coefficients and coefficient p-values for the correlation between total power and wind speed for sites in the EDRP-EDF dataset

The relation between power, wind speed, and precipitation rate is illustrated in Figure 6.5 which shows power-weather distributions for a typical site.

6.1.3 Dependence of power on precipitation

The dependence of total power demand on the precipitation rate was investigated using the same approach as for wind speed. The results are shown in Figure 6.6 and Figure 6.7. The results are very similar to those for wind speed, with the total power dependence being not significant either for the whole site time series or for the low solar gain sample (median -0.2 in both cases). The pair-wise plot in Figure 6.5 of a typical site further illustrates this finding.

It may be intuited that days with low solar gain would be strongly correlated with precipitation,

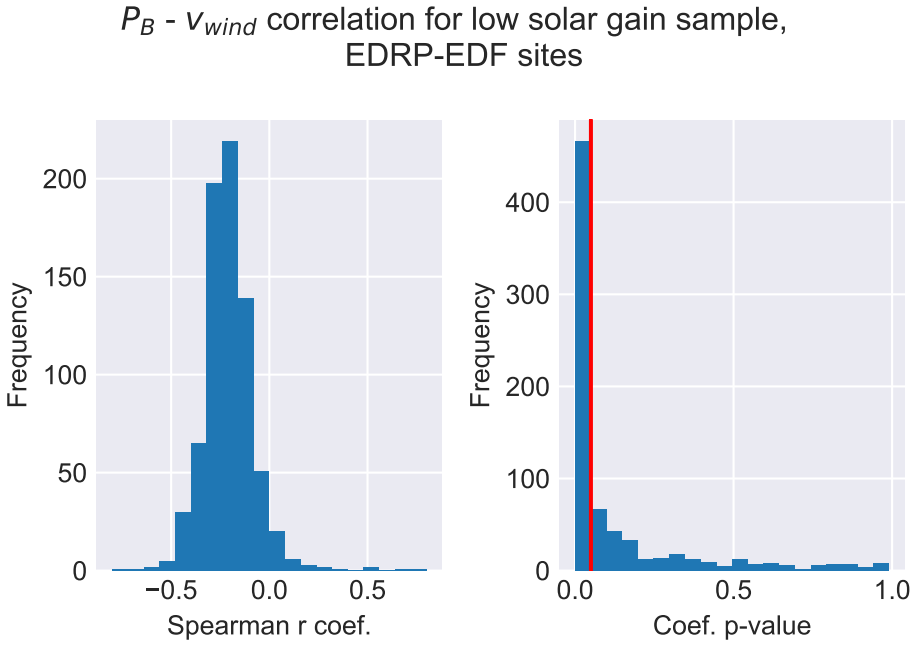

Figure 6.4: Distribution of Spearman rank-correlation coefficients and coefficient p-values for the correlation between total power and wind speed for low solar gain samples of site data from the EDRP-EDF dataset

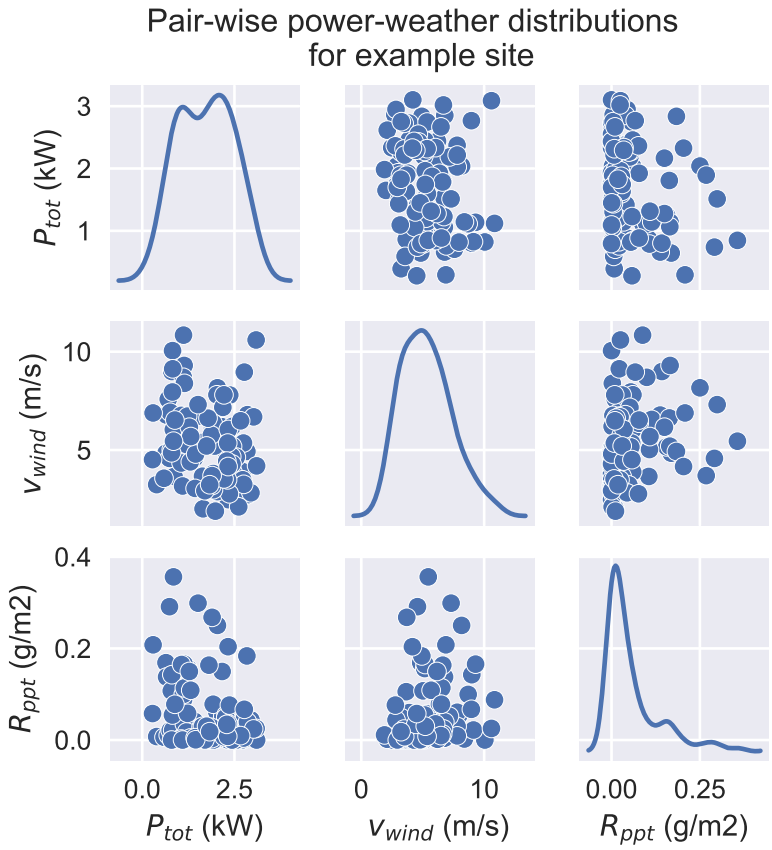

Figure 6.5: Pair wise plot showing the relation between total power and wind and precipitation, with the parameter distribution (smoothed using kernel density estimate) on the diagonal

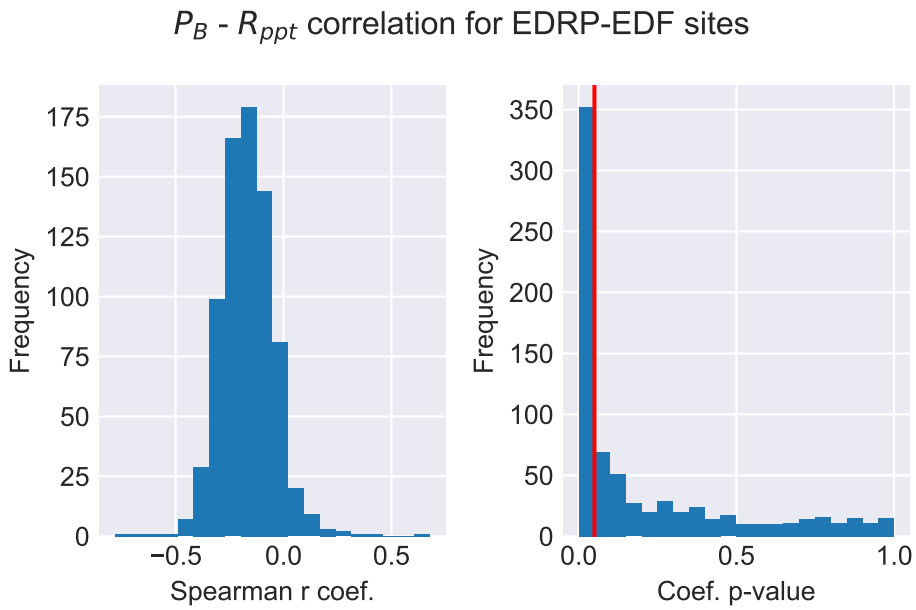

Figure 6.6: Distribution of Spearman rank-correlation coefficients and coefficient p-values for the correlation between total power and precipitation rate R_{ppt} (g/m2 per hour) for sites in the EDRP-EDF dataset

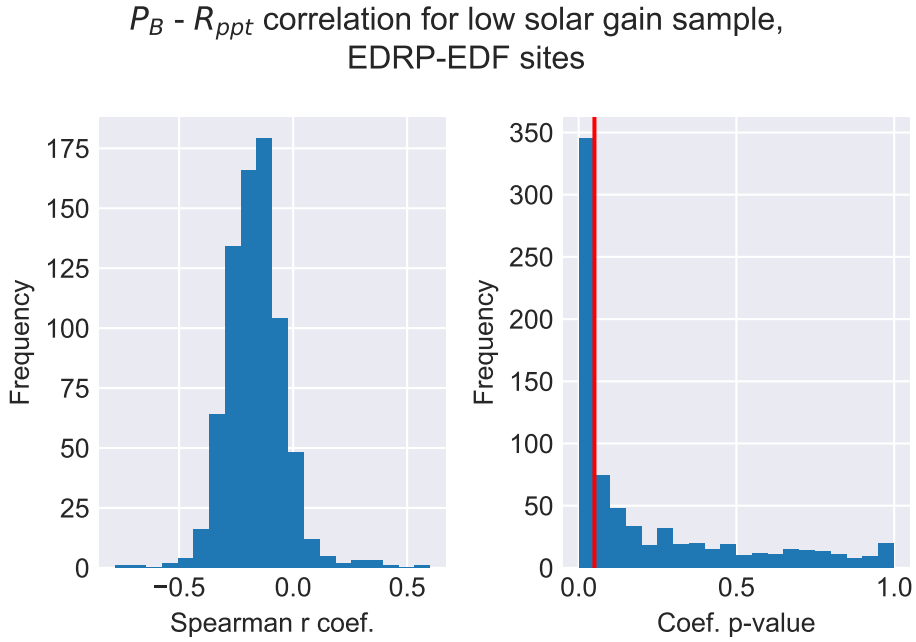

Figure 6.7: Distribution of Spearman rank-correlation coefficients and coefficient p-values for the correlation between total power and precipitation rate (g/m2 per hour) for low solar gain samples of site data from the EDRP-EDF dataset

as they may correspond to days with heavy cloud cover and therefore high chance of rain. This should not have a significant impact on the Deconstruct parameter estimates as it was shown that the power does not appear to depend on precipitation. To verify this, the correlation between solar gains and precipitation for all data was calculated and is shown in Figure 6.8, demonstrating a median correlation of -0.4. This indicates that there is a low correlation between solar gain and precipitation. Combined with the lack of power correlation, it appears reasonable to not consider precipitation in the thermal model.

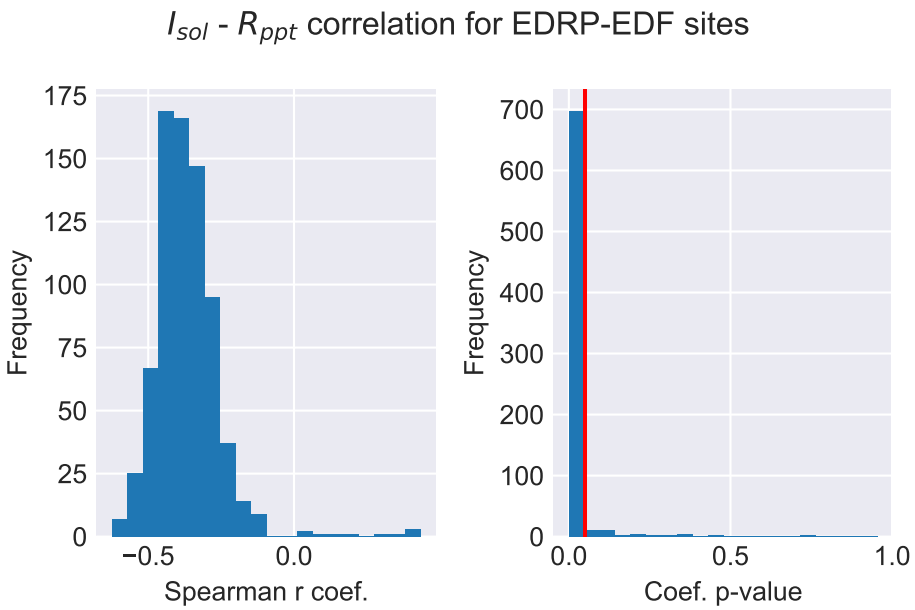

Figure 6.8: Distribution of Spearman rank-correlation coefficients and coefficient p-values for the correlation between solar gains and precipitation rate (g/m2 per hour) for site data from the EDRP-EDF dataset

To further demonstrate that wind speed and precipitation may be safely ignored, the Power Temperature Gradient (PTG) ($PTG = -\frac{dP_{tot}}{dT_{ex}} = -HLC(1 - C_T)$) was calculated using wind speed and precipitation cut-offs. A wind speed upper cut-off of 4m/s was applied based on observations in Everett (1985) which indicate that ventilation exchange is wind dominated above 4m/s. A precipitation rate filter of 0.03 g/m^2 was applied, this threshold was chosen empirically from the precipitation distribution. The distribution of gradients shown in Figure 6.9 demonstrates that there is no change from applying these filters. This is confirmed with a K-S test comparing each distribution against the non-filtered gradient distribution giving a p-value of 0.6 in both cases (i.e. the null hypothesis that the two distributions are the same cannot be rejected).

6.1.4 Evaluating temperature cut-off

In order to evaluate the choice of 15 ° C as an upper temperature cut-off, the PTG ($-\frac{dP_{tot}}{dT_{ex}}$) was calculated for a range of T_{ex} cut-off values from 8 ° C to 22 ° C. Values over 20 ° C should have effectively no effect on the PTG while values below 10 ° C would cut out significant fractions of

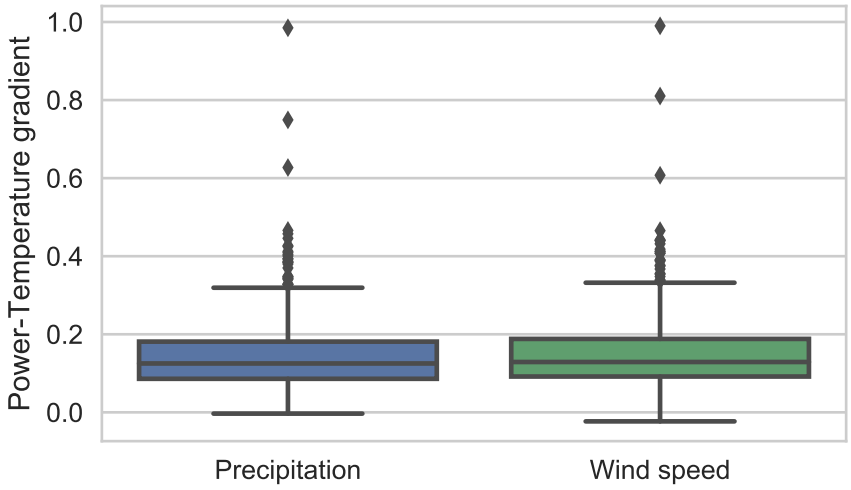

Figure 6.9: Boxplots of power-temperature gradients for EDF dataset with precipitations and wind filters

data. The resulting PTG values and standard errors shown in Figure 6.10 demonstrate that the mean PTG values change rapidly as the cut-off approaches 15 ˚ C but much more slowly above this value, while the standard error drops rapidly from 8 ˚ C to 15 ˚ C and also levels off above this value. From observation of these results, cut-off values in the range of 15-17 ˚ C would be reasonable. 15 ˚ C was selected as the cut-off on the strength of these observations in addition to evidence presented in existing literature (Ruch and Claridge 1992; Summerfield, Oreszczyn, and Hong 2013; Hamilton et al. 2016a; Oreszczyn et al. 2006).

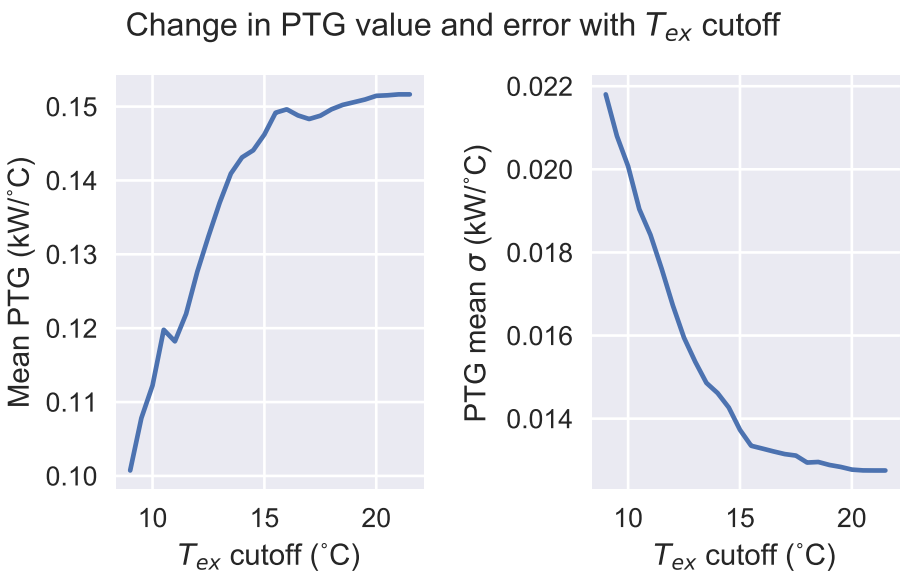

Figure 6.10: Change in the power-temperature gradient estimate (PTG) and the standard error of the PTG as a function of the choice of temperature cut-off value.

6.1.5 Evaluating low solar gain

The impact of the low solar gain filter was evaluated using the EDRP-EDF dataset. Figure 6.11 gives an example of the low solar gain sample resulting from the application of the low solar gain filters to the source data (shown in Figure 6.1).

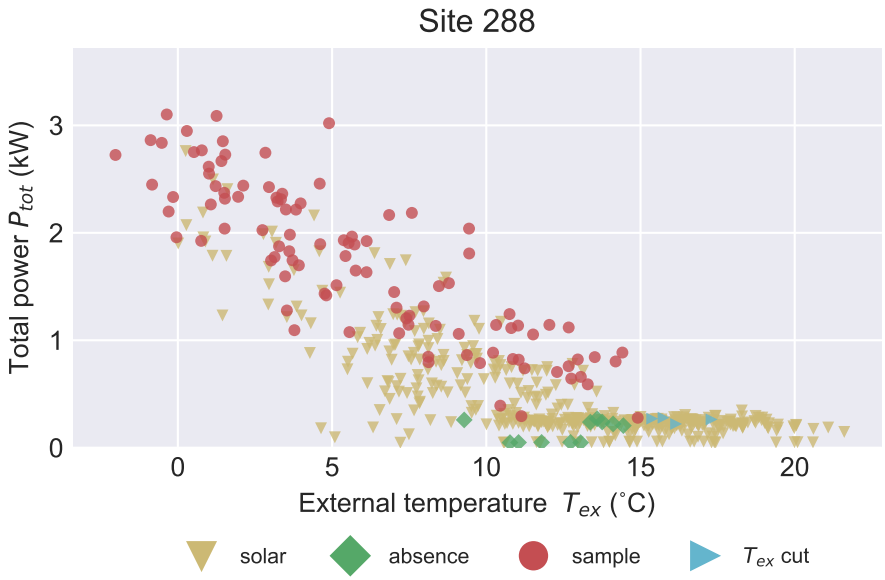

Figure 6.11: Example of low solar gain sample daily average power-external temperature values from EDRP-EDF site 288

In order to evaluate the impact of the data sampling strategy, it was decided to minimise the introduction of additional assumptions by calculating the power-temperature gradient $PTG = -\frac{dP_{tot}}{dT_{ex}} = HLC(1 - C_T)$ only, thereby setting aside temporarily the internal temperature model and the heating system efficiency.

The PTG was estimated using a base filter consisting of a temperature threshold where $T_{ex} < 15\,^{\circ}C$, outlier and absence removal, and was compared to a PTG estimated using the Deconstruct data sampling approach, applying a low solar gain filter before the temperature, outlier, and absence removal filters.

One of the benefits of the solar filter is to reduce the impact of confounding factors. This can be measured be comparing the CVRMSE (root mean square error expressed as a percentage) between the base filter and the low gains filter. It is expected that filtering for solar gains would improve the model fit, resulting in a corresponding decrease in errors.

The number of sites that did not produce a fit was recorded. These are sites where data appeared corrupted or there was not enough data to produce a fit after filtering had been applied. Of the 780 original sites, 763 produced results with the base filter and 743 with the low solar gain filter. When data is filtered from sites with few initial data points there may no longer be enough data for a good fit.

The CVRMSE distributions shown in Figure 6.12 demonstrate a significant reduction through filtering for solar gains. Following recommendations from ASHRAE (2002) a 30% CVRMSE cut-off was applied, above which the model was not considered to be adequately calibrated to the site. 501 sites passed this test with the base filter compared to 614 with the solar gain filter. This represents a 79% overall pass rate with the solar gains filter compared to a 64% pass rate with the base filter. It is interesting to note that in both cases the distributions display a secondary 'bump' of high CVRMSE sites - these may correspond to dwellings with different heating patterns where the simple model currently used fundamentally does not apply.

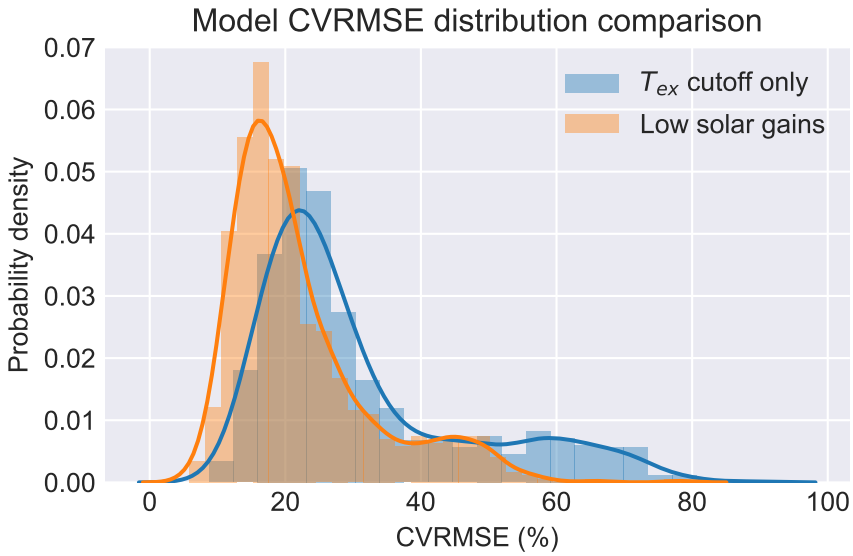

Figure 6.12: Comparison of the distribution of model fit CVRMSE when using a base temperature cut-off only filter and when using a low solar gain filter.

The importance of the low solar gain filter can also be shown through the effect of solar gains on the PTG estimate. In the case that the low solar gain filter has no effect, PTG estimates would be identical, which is not the case as can be seen from Figure 6.13. The distribution of percentage change in PTG for the low gain filter relative to base can be seen in Figure 6.14. The median bias is 13% relative to the low gain sample estimate, i.e. failing to account for, and limit, the impact of solar gains would lead to an overestimate if the dwelling heat loss rate. From inspection of dwelling data, this bias is caused by the tendency of warmer days to also have larger solar gains (Figure 4.4) - the correlation is moderate (median 0.5) but has an impact due to the specifics of the data distribution. Without the solar gain filter, points at higher temperatures will have lower mean metered power demand while points at low temperatures are unaffected, resulting in the gradient appearing larger than when solar gains are limited.

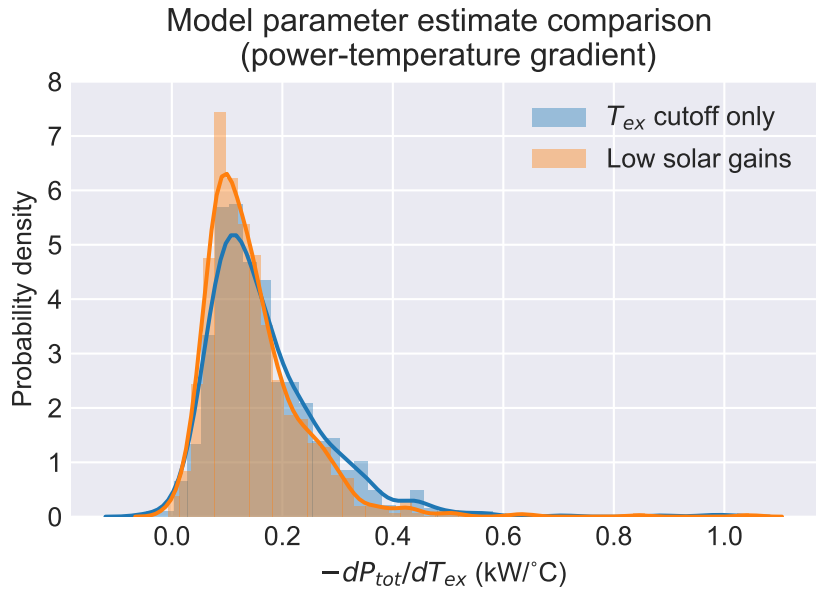

Figure 6.13: Distribution of power-temperature gradient ($-\frac{dP_{tot}}{dT_{ex}}$) when using a base temperature cut-off only filter and when using a low solar gain filter.

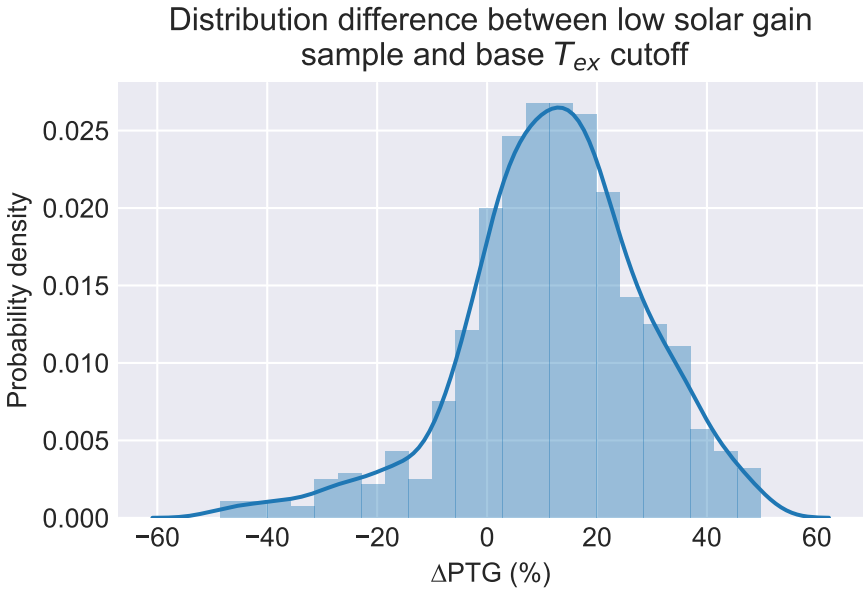

Figure 6.14: Distribution of percentage difference in the power-temperature gradient when using a base temperature cut-off only filter and when using a low solar gain filter.

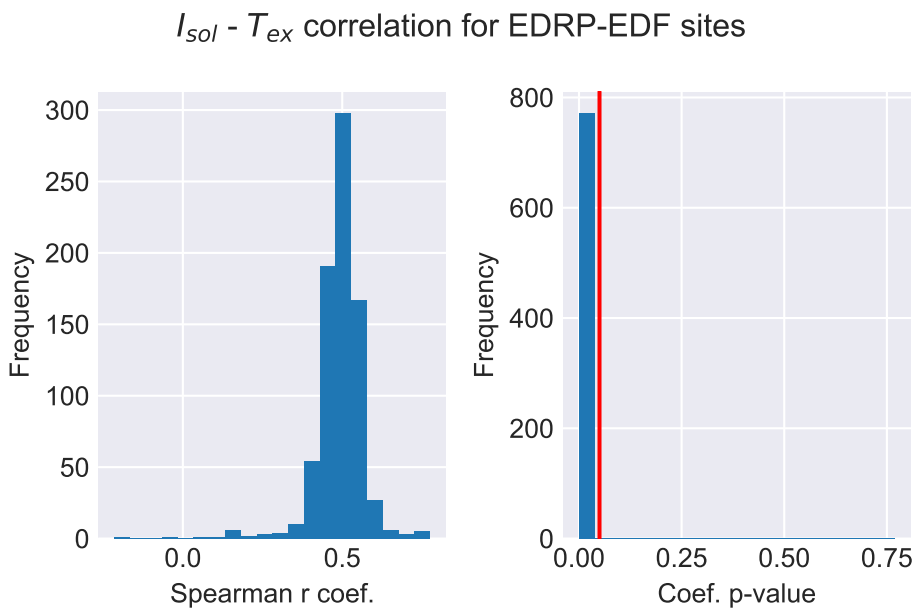

Figure 6.15: Distribution of Spearman rank-correlation coefficients and coefficient p-values for the correlation between solar gains and external temperature for site data from the EDRP-EDF dataset

6.1.6 Model success rate

The EDRP-EDF dataset was used to estimate typical success rates for the model, in terms of the number of dwellings for which the simple thermal model is a reasonable approximation and for which the results meet minimal quality criteria.

As described in Section 4.2, only the sites were retained which had one year of data with less than 50% missing values with both electric and gas meter data. Of the 780 sites with both gas and electric meter data, 654 sites met the data quantity requirement. Applying the model to this subset, one site failed due to the optimiser not finding a solution, resulting in 653 sites having inferred thermal parameters. The model CVRMSE must be under 30%. 542 sites passed this requirement, an 83% pass rate.

The HLC estimates need to have reasonably low uncertainties in order to be useful. As shown in Section 6.3 most estimates have true uncertainties around 15% when the regression covariance is under 30%. By applying a cut-off of 30% an indication of the proportion of sites with acceptable uncertainties. Applying this cut-off in addition to the CVRMSE cut-off leaves 537 sites, an acceptance rate of 82%.

6.2 Internal temperature model

In Chapter 3 Section 3.4 it was demonstrated that the HLC is a function of ΔT, and therefore depends on $C_T = dT_{in}/dT_{ex}$. A linear model of internal temperatures was presented (eq. 3.4), and combined with the thermal model in Chapter 4, Section 3.11.1. This enables modelling internal temperatures in the common case that they are not measured for a dwelling. The temperature model is indexed against a constant T_{fix} of 5 $^\circ$C (278.15K) outdoor temperature, following a similar approach to Hamilton et al. (2016a) and Oreszczyn et al. (2006), such that the internal reference temperature T_0 estimated for the model is the mean internal temperature when the external temperature is 5 $^\circ$C.

6.2.1 Internal temperature covariance parameters from EFUS

In this section, estimation of the covariance parameter C_T of eq. 3.4 is examined. Since the internal temperature is collinear with the total power, it is not possible to estimate C_T from power-temperature data alone. Therefore, a representative covariance parameter for England was estimated using the English Housing Survey (EHS) and English Follow Up Survey (EFUS) data. EHS includes a Government Office Region (GOR) identifier as well as weightings for each site to enable population-representative distributions to be generated. This enabled a representative model for England to be derived empirically.

Daily average internal and external dwelling temperature data was imported from the EFUS dataset and linked with a GOR identifier included in the EHS dataset. Average weather for each GOR was extracted from MetOffice gridded files using the method described in Chapter 5. An example of the internal-external temperature relation is shown in Figure 6.16.

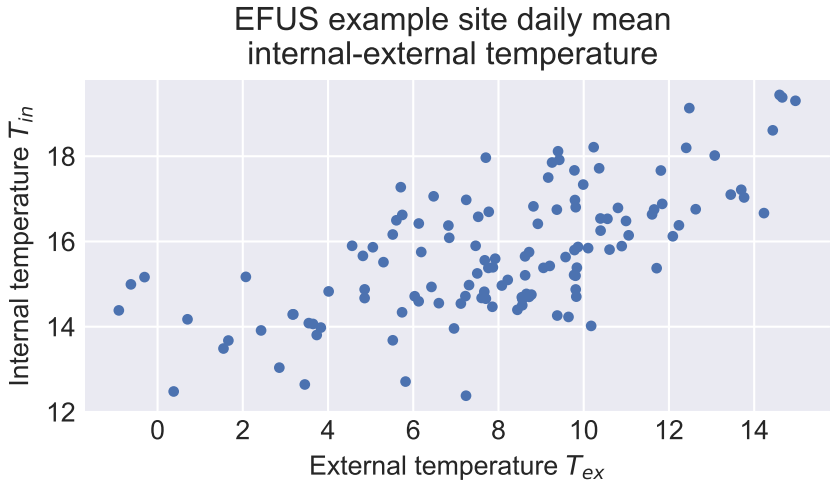

Figure 6.16: Relation between internal and external temperature for a randomly selected site from the EFUS dataset, demonstrating the approximately linear relation between internal and external daily mean temperatures.

The irradiance data associated with EFUS sites was used to produce site data subsets under low solar gain conditions matching those used for estimating HLC. C_T was estimated by performing a linear regression of the mean internal temperature against the external temperature from the EFUS data rather than from the MetOffice linked data. The EFUS temperature data was derived from MetOffice Integrated Data Archive System (MIDAS) observations by associating each site with the nearest weather station. Since this was done before the site location data was removed for archival, it could be assumed that the temperatures would be more accurate than those from the MetOffice. However, the MetOffice data was necessary for the other weather variables, notably irradiance.

The distribution of the temperature covariance parameter C_T for EFUS sites is shown in Figure 6.17.

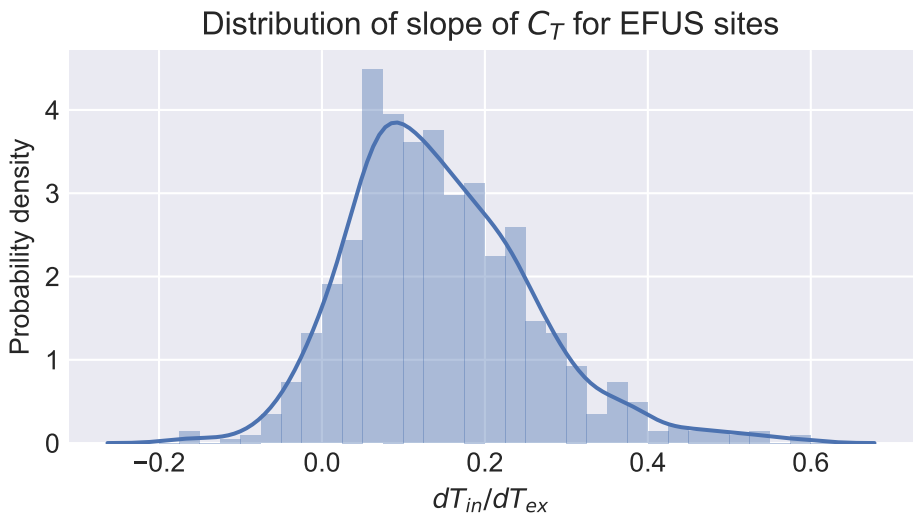

Figure 6.17: Distribution of C_T using EFUS measured internal temperatures

To verify that the linear model was a reasonable approximation of the internal/external temperature dependence, the model fit condition numbers were compared between the linear model and a quadratic model. The condition number can be used as a diagnosis for multicollinearity of the model fit parameters. In general, a higher condition number for a model with additional degrees of freedom indicates that the model parameters are more collinear - in other words, the additional parameters do not contribute additional information. Figure 6.18 demonstrates that the quadratic model had a much higher condition number on all sites, indicating that the linear and quadratic are collinear, and therefore that the introduction of the quadratic term does not improve the model fit.

To obtain a representative mean covariance value for England, a weighted average of the positive gradient values was taken, using site weightings provided by EHS. The resulting mean slope was 0.17. Only positive values were taken since a negative correlation between internal and external suggest a non-thermodynamically driven process. Values that violate the physical assumptions

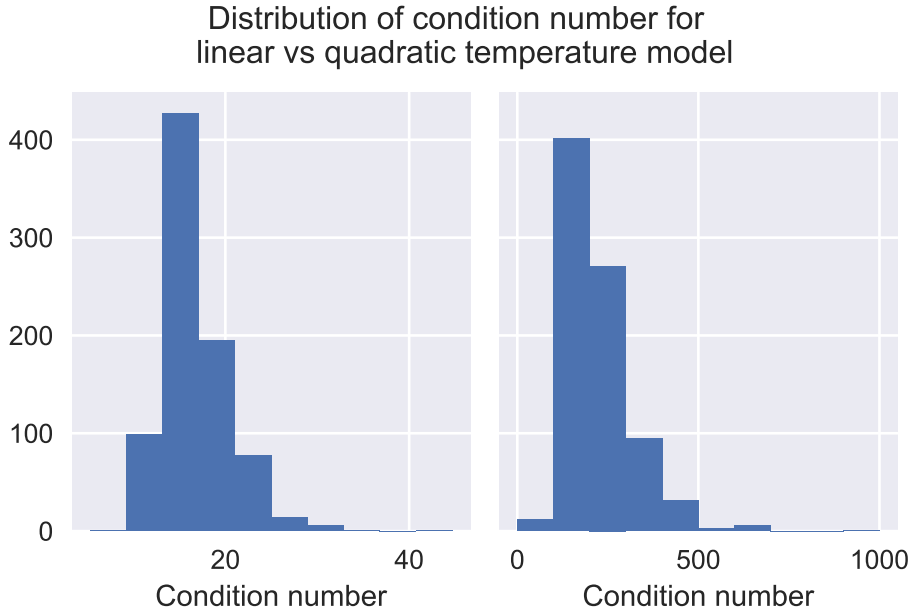

Figure 6.18: Distribution of condition number for linear (left) vs quadratic temperature model (right)

behind the dwelling thermal model were omitted – resulting in removal of removed 6% of weighted sites (Figure 6.19).

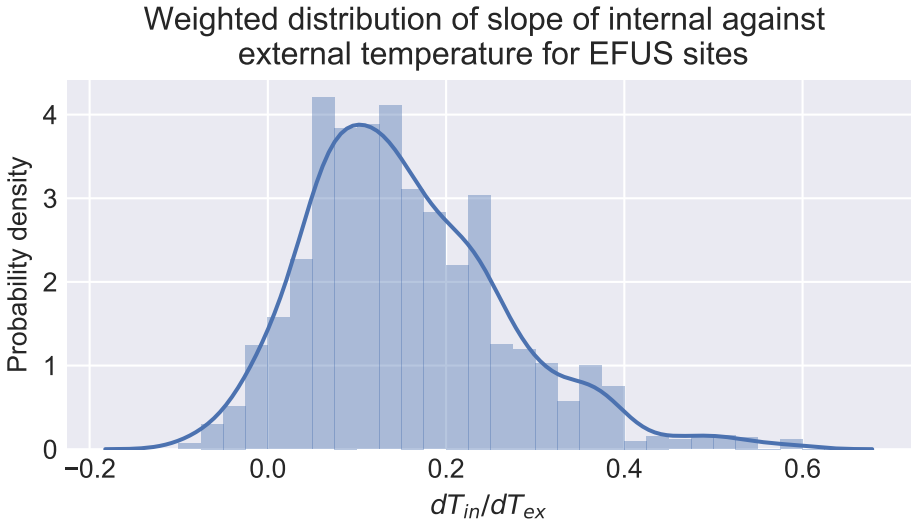

Figure 6.19: Distribution of C_T for EFUS sites weighted by EHS-provided sample weights to obtain nationally representative distribution

Next, the principal physical determinants of the internal temperature covariance were investigated, following the work of Hamilton et al. (2016a). The authors showed that dwelling age and dwelling type were the principal determinants of mean internal temperature. For dwelling type, the authors found that the principal determinant was whether the dwelling was a flat or a house. Comparing the slopes distribution for these two categories reveals no such dependence (Figure 6.20).

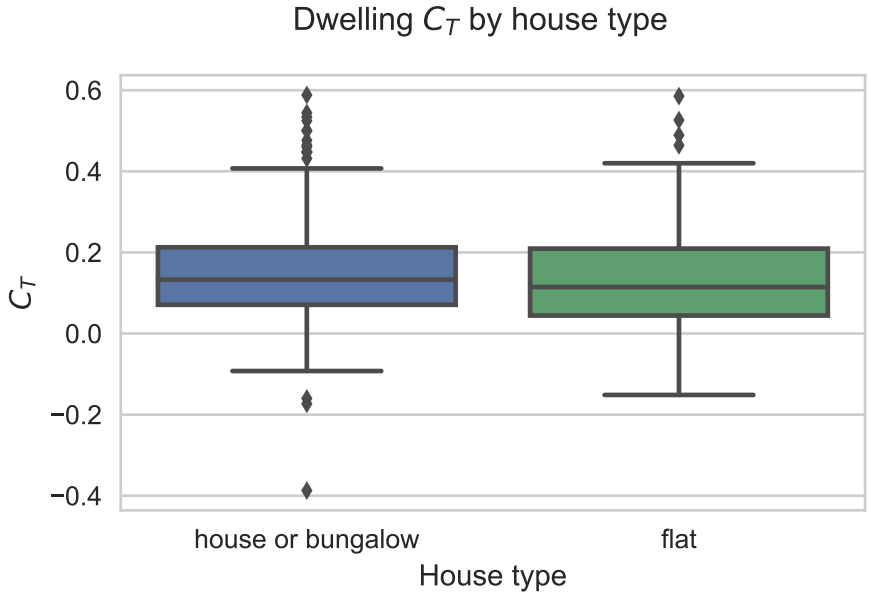

Figure 6.20: Boxplots of covariance between dwelling internal and external temperature for house/bungalow type dwellings compared to flats

For dwelling age, EFUS provides a set of nine age bands (Figure 6.21) as well as simplified 6, 5, and 4 band groupings. As can be seen from Figure 6.21, it is clear that there is a change in the mean slope between oldest and newest dwelling. However the age bands being relatively narrow results in the difference between each successive band not being significant. This was demonstrated using Analysis of variance (ANOVA) tests of band sets. It was found that there was no significant difference between the lower three bands 1850-1918 (p=0.2) nor the upper 4 bands 1965-present (p=1). As a result, the simplified 4 band structure offered by EFUS was used (Figure 6.22).

The results of regression of the covariance against dwelling age are presented in Table 6.1. There is a small (-0.012) but statistically significant (p=0.0005) trend in the covariance. This may be expected for a number of reasons. The improved insulation of newer dwellings would reduce the rate of temperature drop during non-heated periods, thereby reducing C_T. This becomes clear in post-war dwellings where the first energy efficiency measures were introduced (such as cavity walls). Newer dwellings also tend to be heated throughout, while older dwellings may have more unconditioned spaces (e.g. attics, non-heated rooms). The temperature in unconditioned spaces will vary more closely with external temperature, leading to the dwelling mean internal temperature also varying somewhat more with external temperature (Roberts 2008).

Table 6.1: Internal temperature coefficients for each age band.

Age band	C_T
pre 1919	0.18
1919 to 1944	0.16
1945 to 1964	0.15

Age band	C_T
post 1964	0.14

6.2.2 Model coefficients from the SWI dataset

The SWI dataset includes internal temperature monitoring, enabling a similar analysis to the EFUS dataset - regressions for sites with sufficient data are shown in Figure 6.23. Considerable effort was expended in handling the relatively poor quality temperature monitoring data. Of the 82 original sites, only 40 had enough data of sufficient quality to produce a reliable internal temperature model fit. Model parameters C_T, mean=0.19, and T_0, mean=17.3 $^\circ$C, were recorded. Note that $\overline{C_T} = 0.19$ was slightly larger than the pre-1919 mean derived from EFUS.

6.2.3 Model variant effect on HLC estimate

This section compares the variants of internal temperature models and their impact on the estimated HLC against using measured internal temperatures. From eq. 4.3 the HLC estimate should be scaled by C_T.

The SWI dataset was used to compare the temperature model with measured internal temperatures, using the 40 sites that were found to have acceptable internal temperature data in Section 6.2.2. The HLC was estimated using the following variants of the internal temperature model, with values for C_T and T_0 given in Table 6.2:

1. Constant T_{in}: the mean dwelling internal temperature is assumed not to depend on external temperature.
2. T_{in} as linear function of T_{ex}, using an England-wide mean C_T derived from EFUS
3. T_{in} as linear function of T_{ex}, using a mean C_T for dwellings in the pre 1919 age band derived from EFUS
4. T_{in} as linear function of T_{ex} using C_T estimated for each site
5. T_{in} as linear function of T_{ex} using both C_T and T_0 estimated for each site.
6. HLC estimated from ΔT calculated using measured internal and external temperatures.

In variants 1-4, T_0 was estimated from site power-temperature data during the model fit, while in variant 5 it was taken from the linear internal temperature model fit using measured internal temperatures. In all cases SWI data from June 2010 until June 2011 was used. Note that only the EFUS age-band adjusted C_T results are presented and not the mean C_T of SWI. This is done for clarity as the results were very similar for both variants.

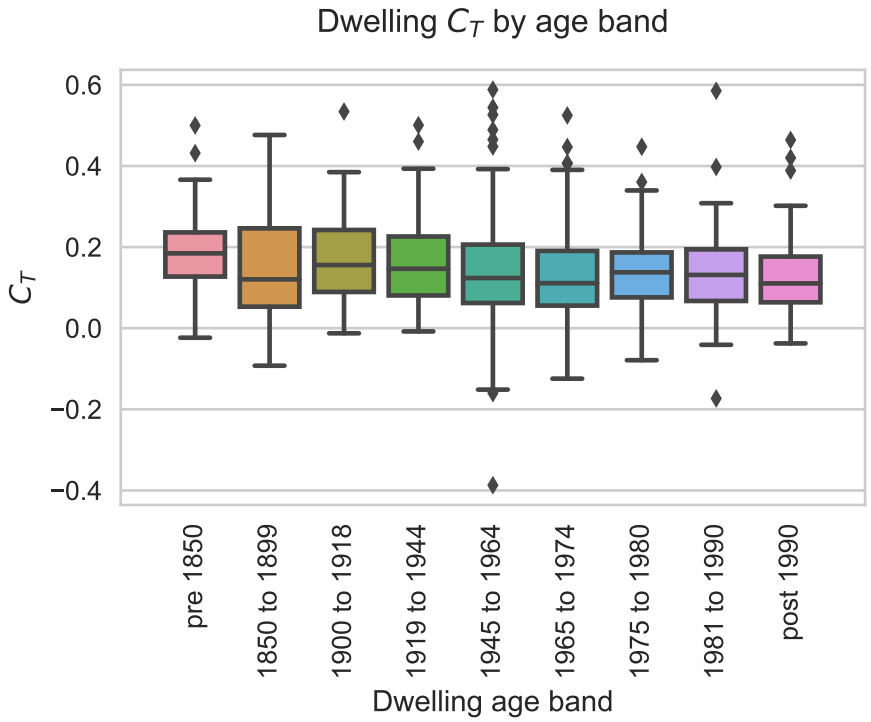

Figure 6.21: Boxplots of covariance between dwelling internal and external temperature for nine age bands

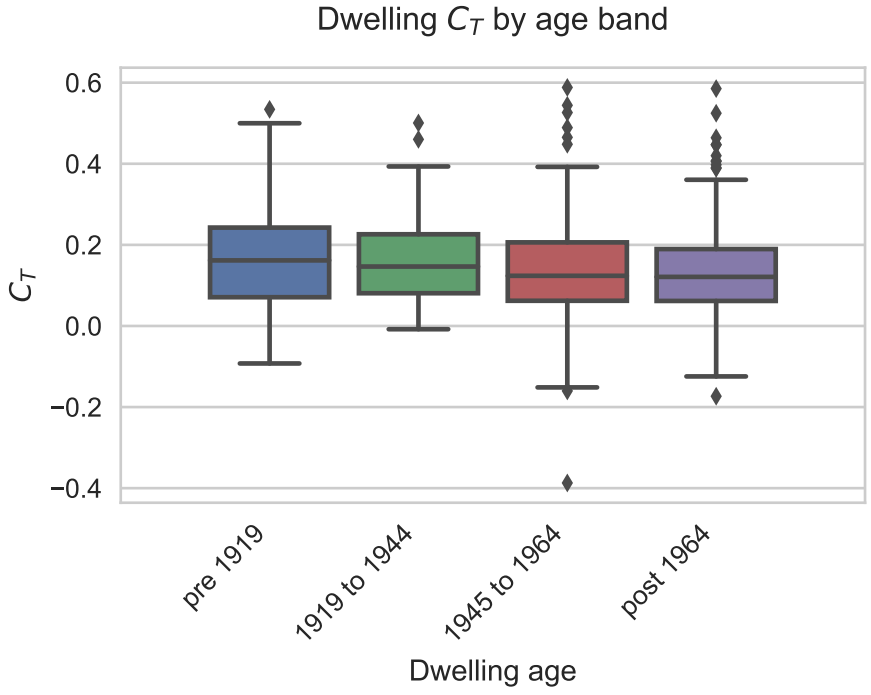

Figure 6.22: Boxplots of covariance between dwelling internal and external temperature for four age bands

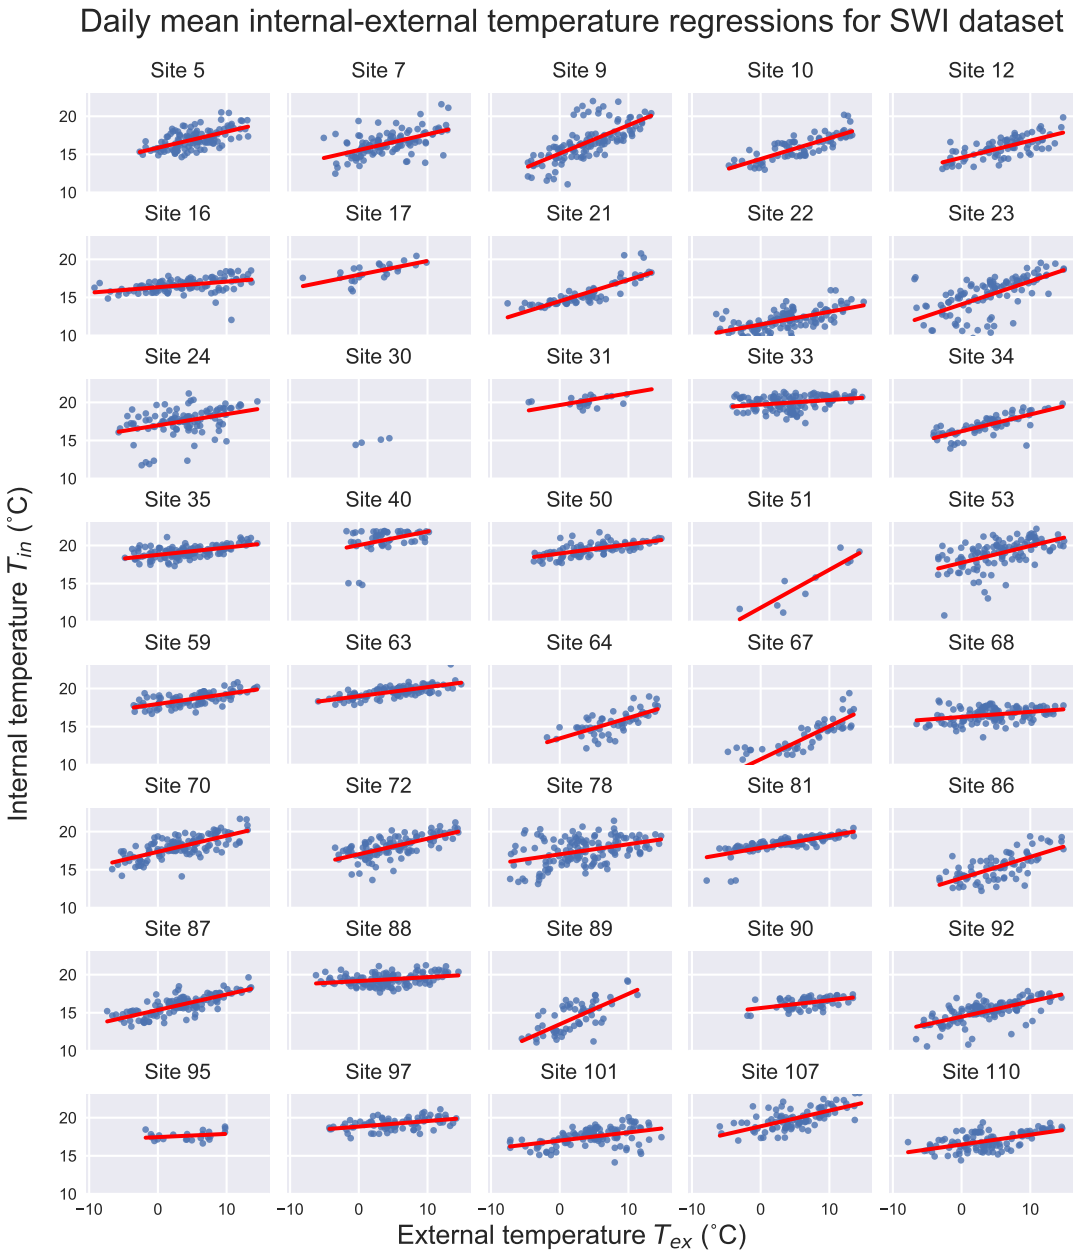

Figure 6.23: Internal-external temperature linear regressions for SWI dataset sites, using a low solar gain sample excluding sites with more than 40% missing data in the winter period.

Table 6.2: Overview of internal temperature model variants

	Variant	C_T	T_0
1	No covariance	0	Estimated from P vs. Tex
2	EFUS mean	0.17	Estimated from P vs. Tex
3	EFUS pre1919	0.18	Estimated from P vs. Tex
4	SWI per site cov.	Estimated per site	Estimated from P vs. Tex
5	SWI full model	Estimated per site	Estimated from measured Tin per site
6	Measured ΔT	N/A	Measured Tin

A low solar gain sample was taken and a simple data quality filter applied that required a minimum of 5 complete days of data and a minimum external temperature range of 6 $^\circ$C, and the HLC was estimated as described in Chapter 4. A filter was applied to remove sites where CVRMSE $>$ 30% and the HLC standard error $>$ 30%. The power-temperature regression lines are shown in Figure 6.24.

The HLC results for all SWI sites for variants 1-5 compared against the HLC calculated using measured ΔT (variant 6) are shown in Figure 6.25. From left to right, model variants make increasing use of site specific measurement data. Variant 1 (left-most plot) makes no assumptions about internal temperature and does not use any data beyond power and external temperature. Variant 5 (right-most plot) uses a per-site linear internal temperature model with both the estimated gradient and reference temperature estimated from measured internal temperatures. As expected the agreement between the measured ΔT approach and the modelled T_{in} approach increases when the model makes use of additional on-site data.

The agreement using the 'SWI full model' is better, even though the measured temperatures are replaced with a simple linear approximation. This is a promising result given the difficulty in reliably measuring internal temperatures. Using a linear approximation enables us to reduce noise and fill gaps in measured internal temperature data.

Using a per site covariance parameter does not appear to significantly reduce the deviation from the 'measured ΔT' variant compared to using a nationwide mean. This is advantageous because being able to use the EFUS mean avoids any internal temperature monitoring requirement. An important observation is that only the full model appears to reproduce the most extreme site HLC values - the other temperature models have a tendency to filter out the highest HLC values.

It can be confirmed that for variants 'EFUS mean', 'EFUS pre1919', and 'SWI per-site cov.', HLC is scaled by C_T, since the power-temperature gradient PTG is constant and only the C_T value changes. Let $PTG = HLC_0$ where HLC_0 is the HLC when $C_T = 0$. Then for a given variant $C_T = 1 - HLC_0/HLC$. For these variants C_T calculated from estimated HLC exactly equals the

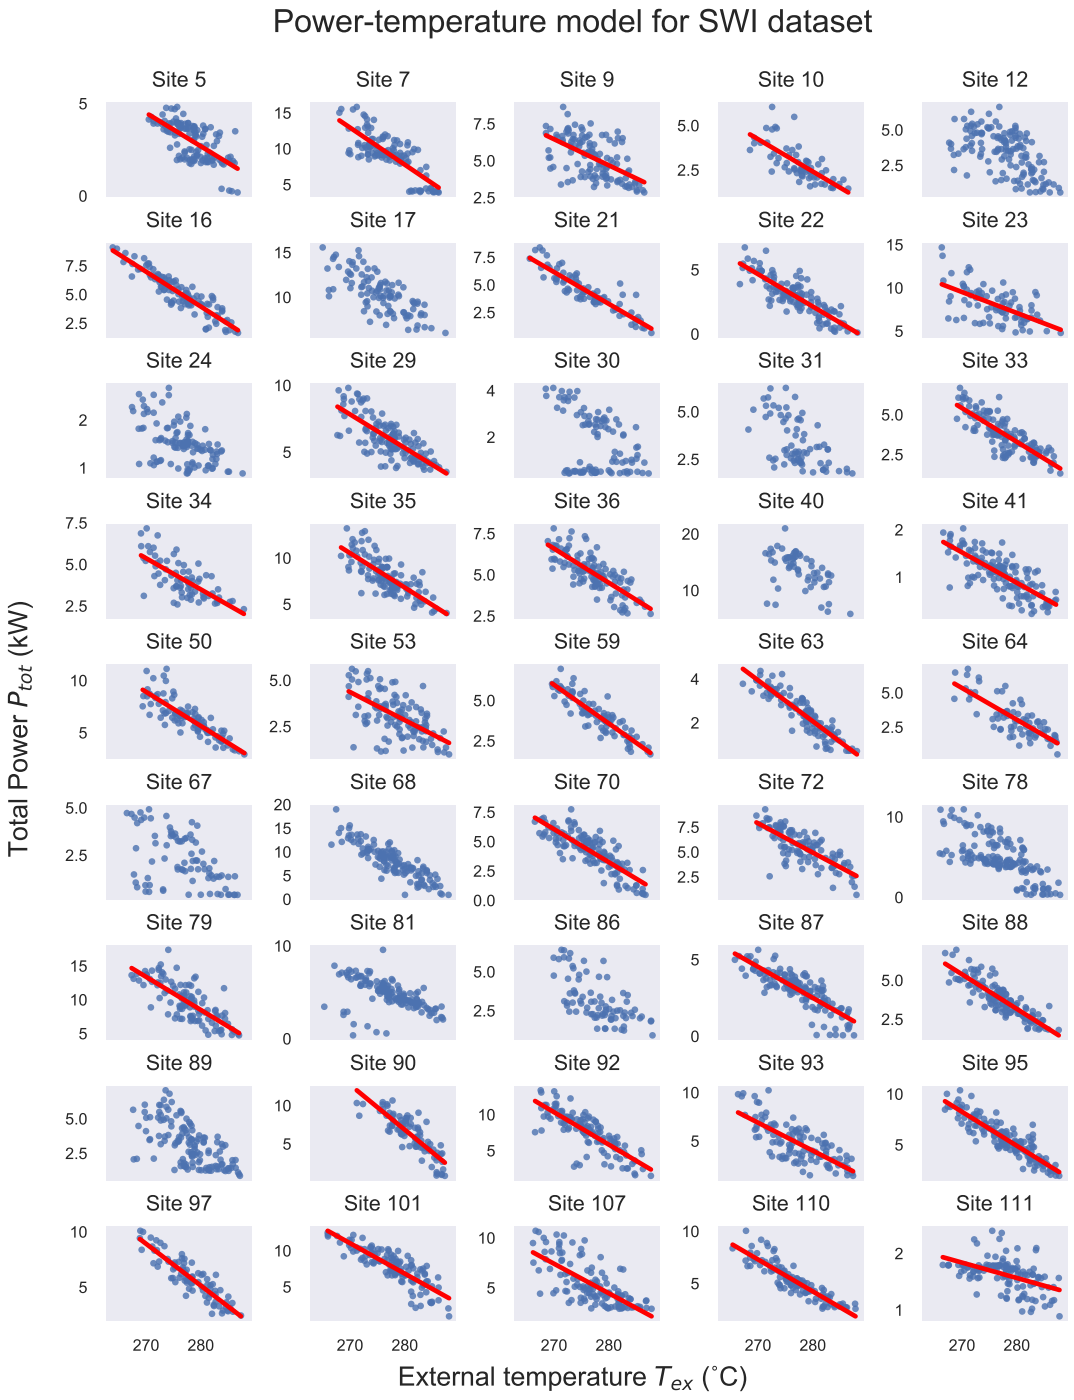

Figure 6.24: Power-temperature models for SWI dataset. Points correspond to the low solar gain sample used for estimating HLC. Red line indicates predicted power using the estimated HLC and reference temperature. Sites missing regression lines are those where CVRMSE > 30% or the standard error on HLC >30%

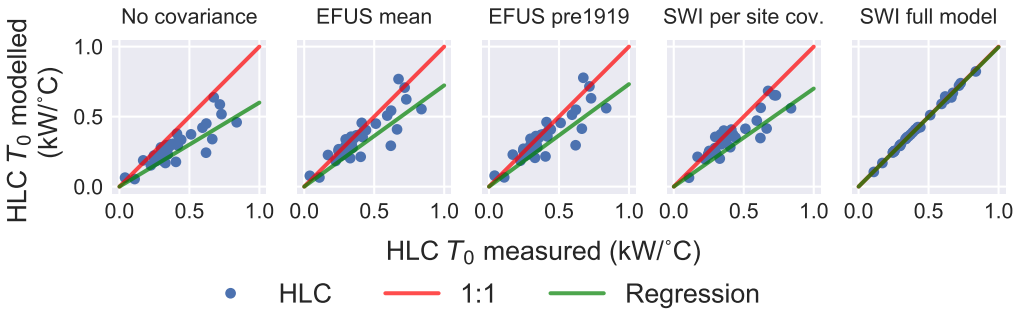

Figure 6.25: Comparison of HLC estimates for SWI sites using the variations on the parameter estimates described in Table 6.2 from left to right: 1: No covariance, 2, EFUS mean, 3, EFUS pre1919, 4, SWI per site cov., 5, SWI full model, against variant 6: measured ΔT

model input C_T. However, this is not the case for the 'SWI full model' variant due to interaction with the externally supplied T_0 parameter.

Figure 6.26 shows distributions of the differences between the ΔT and each model variant. This plot illustrates that there is also some bias in the estimated HLC relative to the ΔT model, also shown in Table 6.3. Bias was significantly reduced by the introduction of the internal temperature covariance parameter, reducing it from 26% to around 10%, while the most significant improvement requires the use of the full internal temperature model, which drastically reduces bias to under 0.2%. The bias adjusted standard error is the standard deviation of the ΔHLC distribution and indicates what the HLC error would be without the bias. This is 15-20% for all variance except for the full model where it drops to under 2%.

Table 6.3: Mean bias, and bias adjusted standard error between HLC estimated with measured internal temperature and using internal temperature model variants.

	Mean bias		Bias adjusted error	
Variant	kW/°C	%	kW/°C	%
No covariance	0.1	26	0.08	15.5
EFUS mean	0.03	9.3	0.06	18.9
EFUS pre1919	0.03	10.0	0.07	19.1
SWI mean	0.02	8.9	0.06	19.3
SWI per site cov.	0.04	11.7	0.06	16.6
SWI full model	0.002 0	.198	0.007	1.62

In order to evaluate the effect of C_T on a larger scale, the analysis was repeated using the EDRP-EDF dataset, which includes dwelling ages but not internal temperatures. The distributions of HLC estimates using mean C_T and dwelling age adjusted C_T are shown in Figure 6.27, revealing that there is effectively no change in this larger dataset.

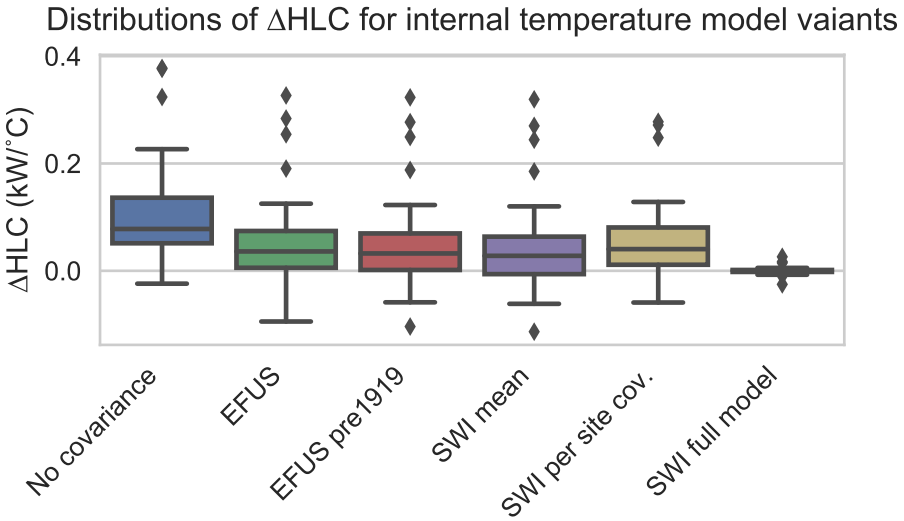

Figure 6.26: Distributions of the difference between HLC estimated with measured internal temperature and HLC estimated using internal temperature model variants.

6.2.4 Temperature model variant effect on T_0 estimate

The internal reference temperature T_0 is the other important parameter inferred by the model. This section evaluates the effect on this parameter of the internal temperature model variants described in the previous section and summarised in Table 6.2.

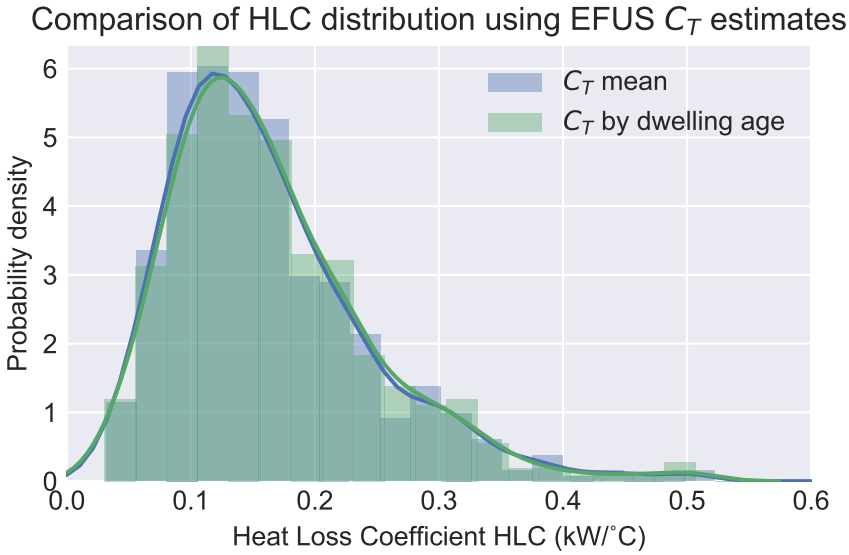

Figure 6.27: Distributions of HLC using EFUS mean C_T and EFUS dwelling age adjusted C_T for HLC estimates for the EDRP-EDF dataset

Plots of the measured mean T_{in} against the modelled T_0 estimate (Figure 6.28) demonstrate that using model variant 5 (using per site C_T and T_0) gave close agreement between measured and modelled internal temperatures. However, the T_0 estimates using other models displayed considerably more scatter than shown when estimating HLC (Figure 6.25), reflecting the larger uncertainty in the T_0 as well as its dependence on assumptions for parameters η_B and η_{HS}. There is a significant increase in scatter between variants 5 and 4, while variants 1-3 are simply scaled relative to variant 4. The estimated T_0 is particularly high relative to the measured value for variant 1, while making use of mean C_T estimates in variants 2 and 3 somewhat reduced the over-estimation.

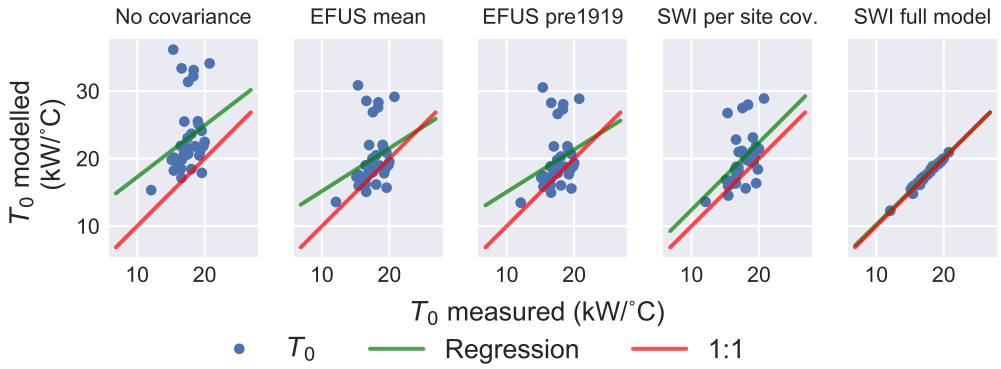

Figure 6.28: Comparison of T_0 estimates across SWI sites for progressively more detailed internal temperature models, see above.

The mean delta and standard deviation of the difference between the estimated and measured T_0 are shown in Table 6.4 and Figure 6.29. These use the T_0 reference temperature from the SWI full model, since the measured temperature mean would not give expected internal temperature at 5 ˚C external as desired for a coherent comparison. These plots demonstrate that using the linear

internal temperature model produces significantly improved internal temperature estimates compared to assuming a constant internal temperature, for all variations of the covariance coefficient estimation. With the internal temperature model, reference temperatures were overestimated by about 1 ˚ C, although there is a large spread of values.

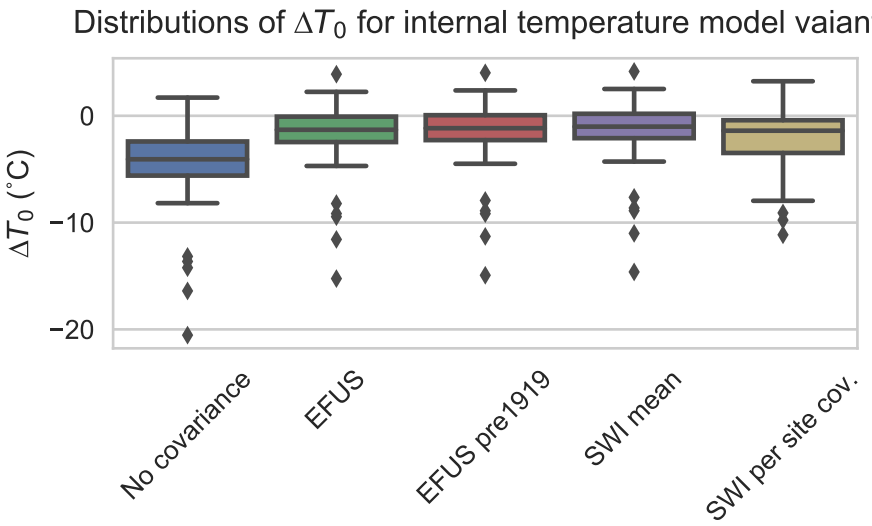

Figure 6.29: Distributions of the difference between T_0 estimated with measured internal temperature and T_0 estimated using internal temperature model variants.

Table 6.4: Bias and mean of the difference between T_{in} estimated with measured internal temperature and HLC estimated using internal temperature model variants

Variant	Mean Δ (˚ C)	Standard deviation (˚ C)
No covariance	5.0	4.1
EFUS mean	1.5	3.1
EFUS pre1919	1.2	3.0
SWI mean	1.0	3.0
SWI per site cov.	1.5	2.1

As in the previous section, the analysis was repeated using the EDRP-EDF dataset and the mean parameter estimates derived from EFUS. The distributions of T_0 estimates using mean C_T and dwelling age adjusted C_T are shown in Figure 6.30. In Figure 6.31 the two estimates are compared against each other. This indicates that there is little advantage to using dwelling age adjusted C_T since the results are internally consistent, however without measured internal temperatures it is not possible to determine how close these results are to reality.

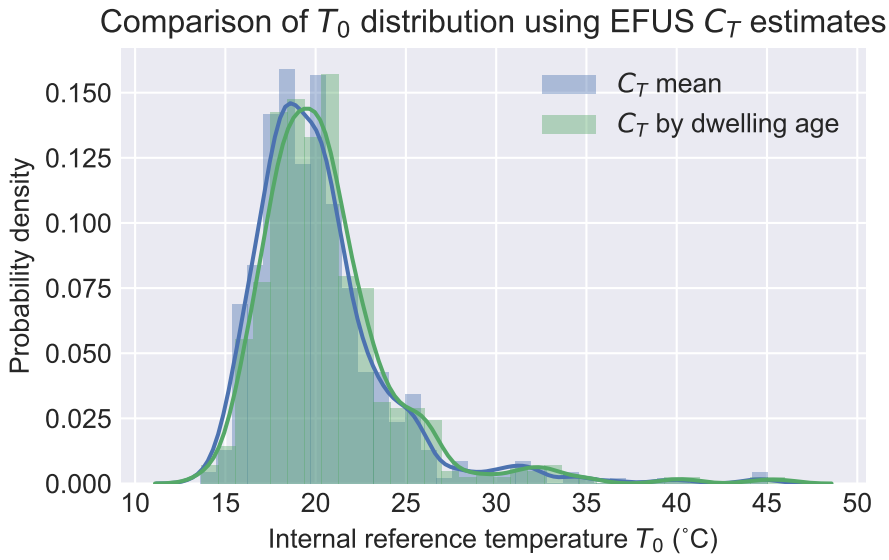

Figure 6.30: Distributions of internal reference temperature T_0 using EFUS mean C_T and EFUS dwelling age adjusted C_T

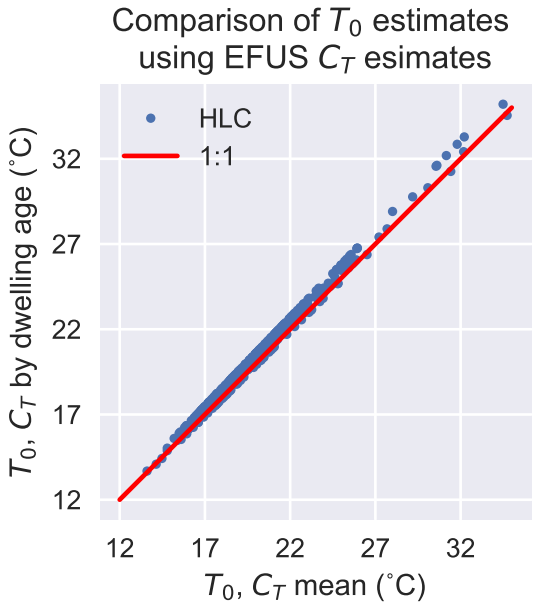

Figure 6.31: Comparison of T_0 estimate using dwelling age adjusted C_T against EFUS England mean C_T

6.3 Model Uncertainty

The following sub-sections provide an estimation of the uncertainty in the model inputs and evaluate their impact on model parameter estimates.

6.3.1 Uncertainty in total power

The uncertainty in the total power is the sum of the uncertainties in the electricity and gas power readings. This research assumes that the error in the electricity meter reading is negligible. The uncertainty in gas power could be significant however because gas consumption is measured through volumetric flow and converted to an energy value. This process introduces uncertainty, notably because gas energy density depends on its temperature and pressure which may vary at the metering point. In this section, mean error in the total power is estimated.

To estimate error in total power, firstly the error deriving from the gas energy content calculation is considered. The UK official formula for the energy content of domestic gas (NationalGrid 2015; Ofgem 2000) used for the calculation is given in eq. 3.15. The Calorific Value (CV) of gas in the distribution lines is regularly tested with high accuracy instruments, and is therefore believed not to be a significant source of error (Lander 2012). The practice is for Ofgem to require that (absolute) error in CV measured by CV determination equipment should not exceed 0.10 MJ/m3. Using a typical mean CV of 38.1MJ/m3 as reference, this corresponds to a 0.3% error in the CV.

The principal source of error in this calculation is obtaining the correct gas volume conversion factor. Ideally, the gas meter would also measure temperature and pressure, in practice this is only common for large industrial consumers (over 2.928 MWh per year) (Ofgem 2000). Ofgem performed a study which found that the average temperature at the gas meter was 12.2 $^\circ$ C (285.35K) while gas pressure is affected by atmospheric pressure, which in the long term depends only on the altitude. A standard mean altitude of 66m above sea level was chosen. Combining these assumptions produced the correction factor 1.02264 to convert from volume measured at the meter under the assumed conditions to volume at standard 15 $^\circ$ C, 101.325 kPa conditions to match the CV.

To determine the total uncertainty, the uncertainty in the volume measurement is calculated. The temperature/pressure conditions at the meter are constantly changing, but addressing this is out of scope of this research. Instead, for simplicity, an estimate of mean uncertainty in the winter period is derived. Ofgem (2000) gives high/low worst-case error estimates, but it is improbable for a given site to have the most extreme errors on all inputs (Table 6.5). These values therefore overestimate the mean error in the results.

Table 6.5: 'Low' and 'High' extremes of gas meter measurement error given by Ofgem for the meter device and meter mean temperature and pressure (Ofgem 2000)

Component	Measurement error 'low' (%)	Measurement error 'high' (%)
Meter	-2.0	+2.0
Temperature	-1.46	+1.32
Pressure	-0.79	+3.05

To estimate the distribution of total volume error, a Monte Carlo sampling method was used following the guidelines of the Joint Committee for Guides in Metrology (Joint Committee for Guides in Metrology (JCGM) 2008). This is a numerically simple method to combine distributions of errors of the gas measurement volume, and was useful because the inputs were not guaranteed to be normally distributed. First, suitable distributions across the population of gas customers were defined for volume errors caused by the meter, temperature, and pressure. Then, N=1,000,000 samples of percentage volume deviation were drawn independently from each error distribution and the total gas volume deviation for each draw was calculated as the sum of the independent volume errors. This produced a new distribution for total error.

For the gas meter error distribution, the legally permitted meter volume measurement error of 2% (European Parliament 2004) was used as the standard deviation of a normal error distribution (with mean=0), giving the Cumulative Distribution Function (CDF) in Figure 6.32.

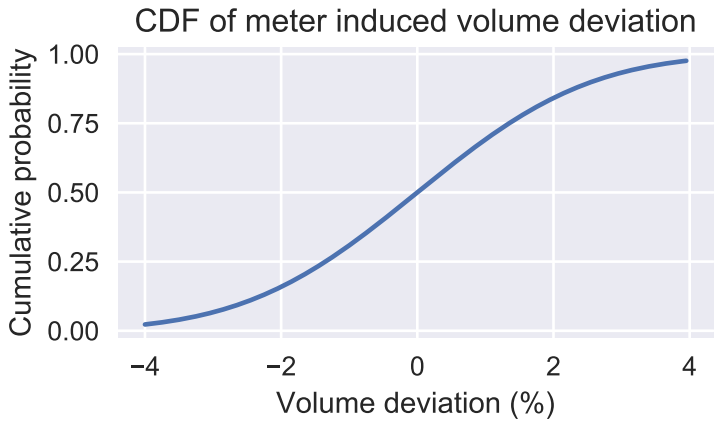

Figure 6.32: Meter volume measurement error CDF, based on a legally permitted error of 2%

Deriving the change in volume of natural gas in the distribution system as a function of temperature and pressure is relatively complicated. Ofgem (2000) gives the following relations between deviation from assumptions for temperature and pressure and error in measured volume:

- 0.35% error per 1 °C deviation from 12.2 °C assumption.

- 0.12% error for each 10m of deviation from 66m altitude assumption.

Based on the Ofgem data in Table 6.6, the distribution of gas customers across altitudes is illustrated in Figure 6.33. For simplicity, the change in volume for each band is approximated using the centre of that band (as opposed to interpolating across the range), giving the CDF shown in Figure 6.34. The Monte-Carlo method is helpful for integrating this distribution since it is clearly non-normal.

Table 6.6: Distribution of domestic consumers by altitude (Ofgem 2000)

Altitude Band (m)	N consumers (millions)	% of consumers	Change in volume (%)
0 to 66	10.21	58	+0.8% to 0
66 to 150	5.87	33	0 to -1
151 to 230	1.33	8	-1 to -2
231 to 320	0.14	1	-2 to -3

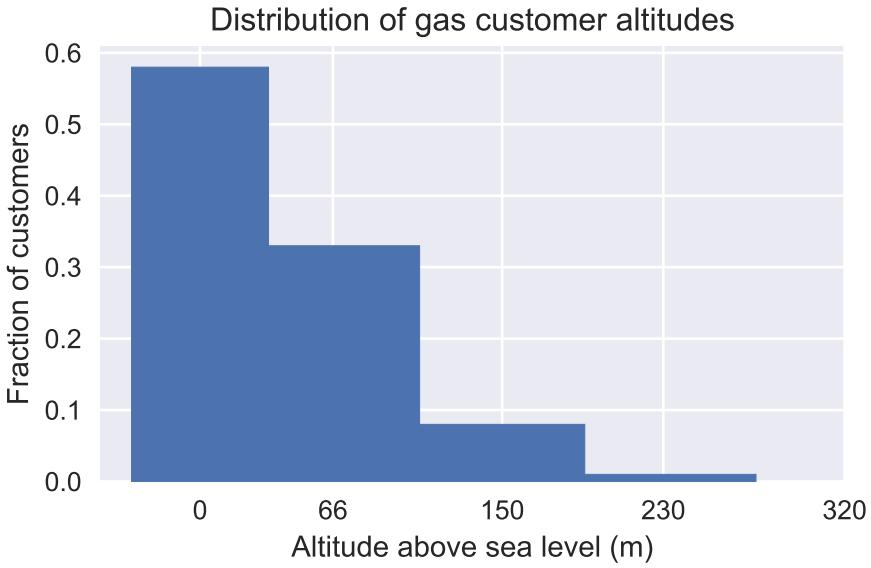

Figure 6.33: Distribution of customers across altitude ranges above sea level, according to Ofgem

Although Ofgem (2000) and Lander (2012) both refer to a study measuring the temperature at gas meters from which the 12.2 °C meter temperature average was obtained, the raw data for this study needed to enable a temperature distribution to be derived was not provided. However, the percentage volume distribution may be approximated using the largest 'high/low' values from Table 6.5 as the standard deviation of a normal distribution (Figure 6.35).

Measurements of external temperature and mean internal temperature from EFUS (DECC 2013b) were used to obtain the distributions in Figure 6.36 by taking the average temperature for each site for the period from 1st October to 31st March, based on the observation that most heating occurs

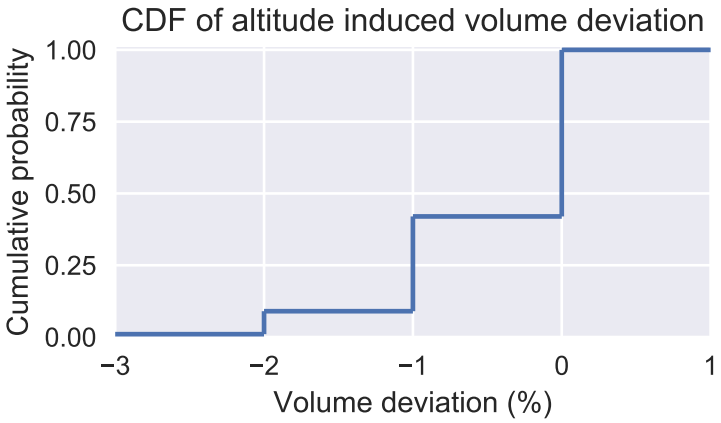

Figure 6.34: Altitude induced volume measurement error CDF, showing three discrete bands for absolute volume deviation.

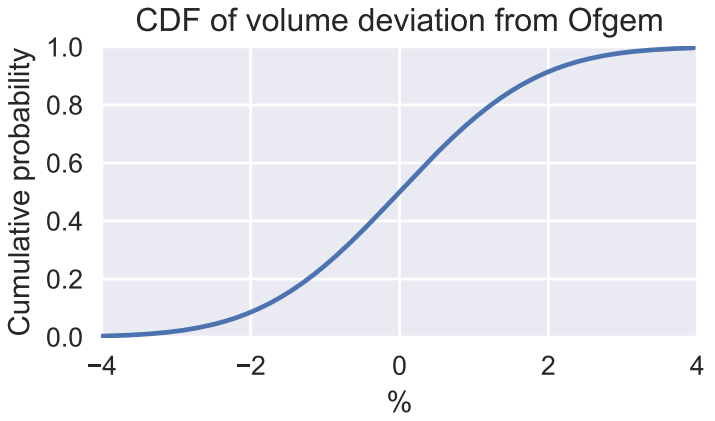

Figure 6.35: CDF of volume change % due to temperature difference at meter from 12.2 ° C assumption, based on std = max(high,low) of Ofgem provided delta values

during this period (Hulme, Beaumont, and Summers 2013a). These were weighted according to the site weightings provided by EFUS by repeating the entries for each site by the corresponding number from the EFUS weighting sheet.

To estimate bounds on the meter temperature, empirical CDFs were generated for the indoor and outdoor temperatures. These were converted to distributions of volume deviation by calculating the temperature delta from the 12.2 ° C and multiplying by 0.35% error per 1 ° C deviation, and distribution CDFs generated Figure 6.37.

Estimating the distribution of absolute volume percentage change from assumption using the Monte Carlo simulation method gives the distribution in Figure 6.38 with a median error of 3%.

The distributions of uncertainties when using the EFUS temperature data as input is shown in Figure 6.39.

The total uncertainty distribution was obtained by taking the magnitude of the volume percentage deviation (Figure 6.40), which gave median errors of 4% in both cases.

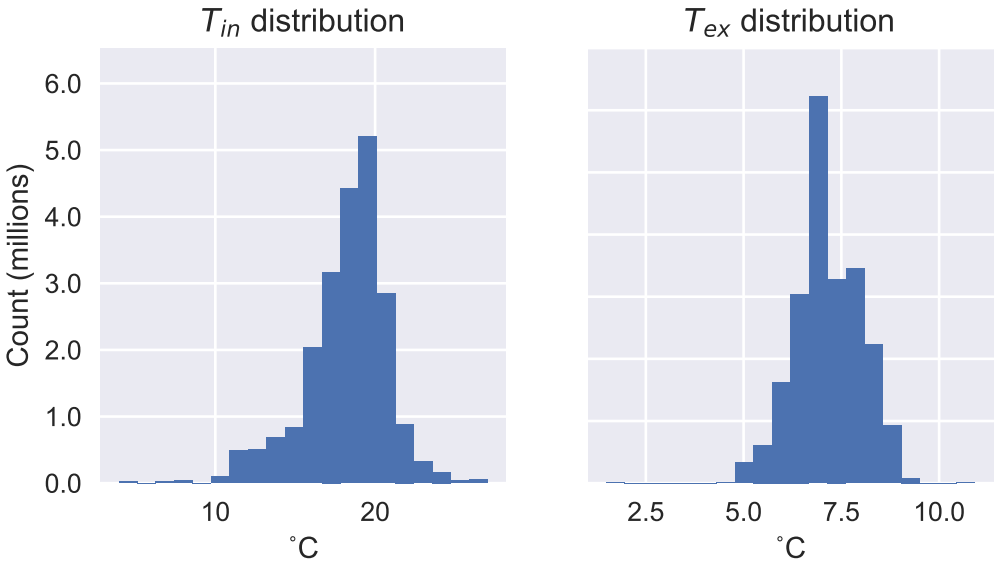

Figure 6.36: Weighted distributions of site internal and external winter temperatures

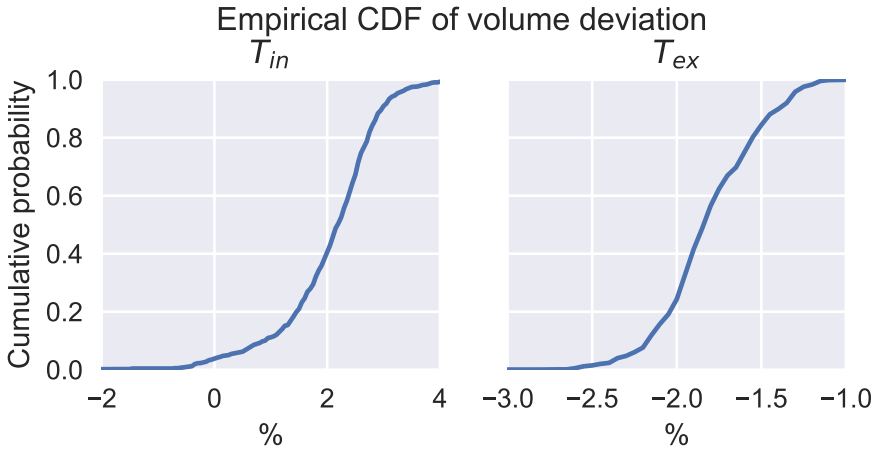

Figure 6.37: Volume deviation percent CDFs resulting from temperature difference between 12.2 ˚ C benchmark and internal temperature (left) and external temperature (right).

The estimated mean error of 3% is smaller than either the worst case estimate provided by Ofgem despite using the same input assumptions, illustrating the importance of considering the joint probability distributions of uncertainties. Using EFUS temperatures as inputs results in similar errors. Therefore the error in gas volume measurement was set to 3%. Combining this with the CV error of 0.3% using the standard root-sum-square of percentage uncertainties approach to uncertainty combination gives (Hogan 2006):

$$\sigma_{P,\%} = \sqrt{(\sigma_{V,\%})^2 + (\sigma_{CV,\%})^2}$$
$$\sigma_{P,\%} = 3\%$$

That is, a total uncertainty in gas power of 3% - as expected, the uncertainty in CV has a negligible

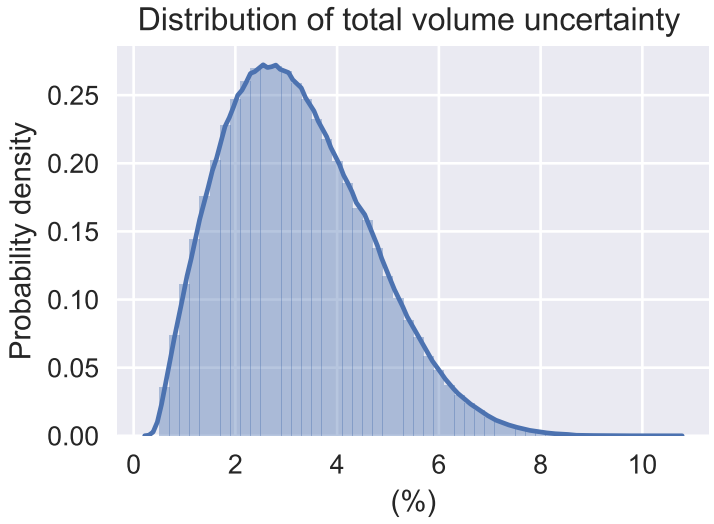

Figure 6.38: Distributions of total absolute % error in gas volume measurement resulting from Monte Carlo method summation of input factor distributions

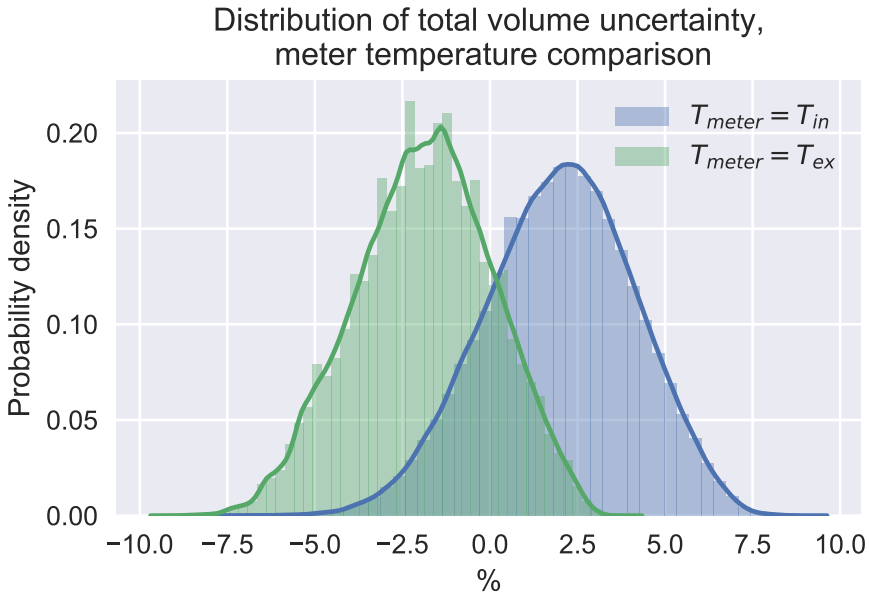

Figure 6.39: Distributions of total % error in gas volume measurement resulting from Monte Carlo method summation of input factor distributions

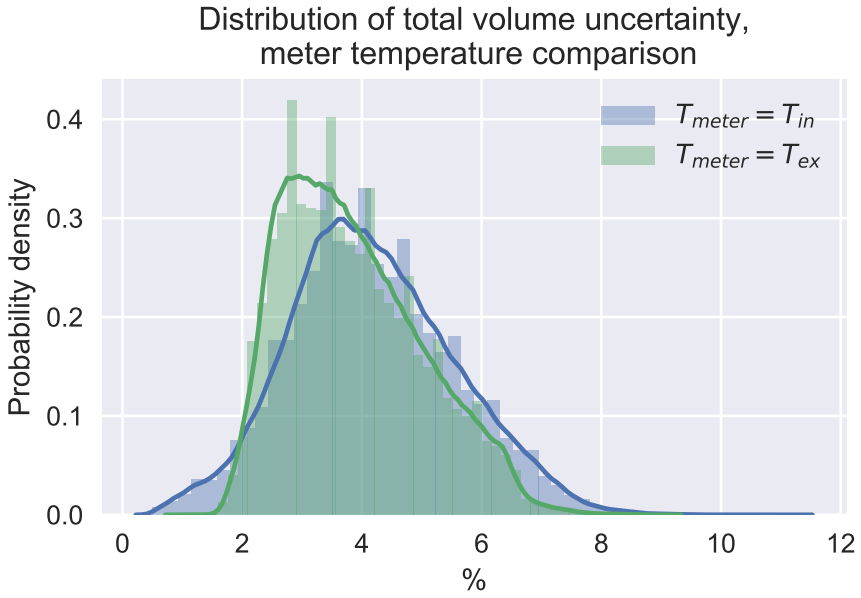

Figure 6.40: Distributions of absolute total % error in gas volume measurement resulting from Monte Carlo method summation of input factor distributions

effect.

6.3.2 Uncertainty in external temperature

Although weather data is provided with a high degree of precision (to 4 decimal places in the case of MetOffice data), this is not representative of the uncertainty in the weather parameters used for energy demand estimates. The weather is drawn from grid squares representing averages over that region. In some cases where site locations are not known exactly averages over several grid squares are used, potentially introducing errors relative to the true weather conditions at the site. Even when weather is measured using on-site weather stations, these may be subject to a range of local effects (shading, wind, urban micro climates) which introduce uncertainty into the readings.

For the external temperature readings, an estimate of the temperature was obtained by comparing EFUS external temperature data (derived from MIDAS) with the MetOffice gridded data. The distribution of difference was found to be normal with a standard deviation of 1 $^\circ$ C. This distribution was incorporated into the Monte Carlo (MC) simulation to estimate the resulting uncertainty in model parameters.

6.3.3 Uncertainty in HLC and T_0

The model parameters are estimated by fitting the thermal model using a standard least squares optimiser. As one of its outputs, the optimiser produces a covariance matrix for the model parameters. A first order estimate of parameter standard error σ can be calculated as the square root of

the matrix diagonal, i.e.:

$$\sigma_i = \sqrt{cov_i i} \tag{6.1}$$

The Monte Carlo (MC) approach described by the Joint Committee for Guides in Metrology (JCGM) (2008) was used to obtain a more robust uncertainty estimate. This propagates uncertainty by using known input uncertainty distributions to generate N variations of the input data perturbed using values drawn from these distributions. The input distributions were derived from the uncertainty CDFs defined in Sections 6.3.1, 6.3.2.

Model parameter estimates were performed using the 780 sites from the EDRP-EDF dataset and the covariance derived standard errors recorded. For the MC model, 1,000 iterations were performed using the following steps:

1. For each run, a value of C_T was selected from the EFUS distribution in order to simulate the possible range of values for a site.
2. For each run, a percent deviation value for the meter and the altitude was randomly generated from the corresponding CDF.
3. For each time step in the power time series, a meter temperature percent deviation was generated.
4. The gas percent volume deviation was calculated and combined with the constant electricity meter uncertainty. A total power deviation was calculated from the percentage and added to the total power time series.
5. For each time-step in the external temperature time series, a temperature absolute value deviation was chosen and added to the time series.
6. The HLC was estimated using the modified time series.

The distribution of covariance absolute and percent standard error and MC standard error for HLC is shown in Figure 6.41. The mean standard error for HLC from the MC approach is 0.02kW/$^\circ$C or 15%. This is equal to the 15% uncertainty estimated by Jack (2015) when performing HTC estimates on simulated occupied dwellings. These results suggest the HLC is robust to perturbations in the input data.

The MC percentage uncertainties are larger than the regression percent standard error but are much more narrowly distributed. Since the regression uncertainties are much faster to compute, it is useful to define cut-offs in terms of these which should be equivalent to acceptable MC uncertainty. Figure 6.42 demonstrates that values with under 30% regression standard error should correspond to sites with under 20% 'true' (MC simulated) uncertainty, with the bulk of sites being between 14 and 16 %.

The distribution of covariance absolute and percent standard error and MC standard error for T_0

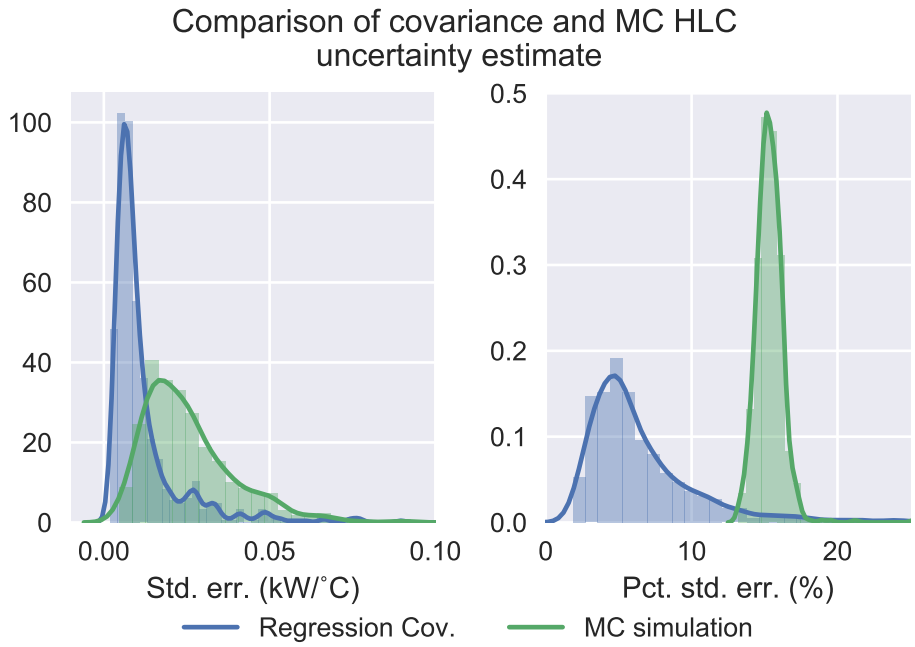

Figure 6.41: Distributions of absolute (left) and percent (right) uncertainty in HLC across 780 EDRP-EDF sites, using uncertainty derived from covariance and using MC propagation of input uncertainty distributions

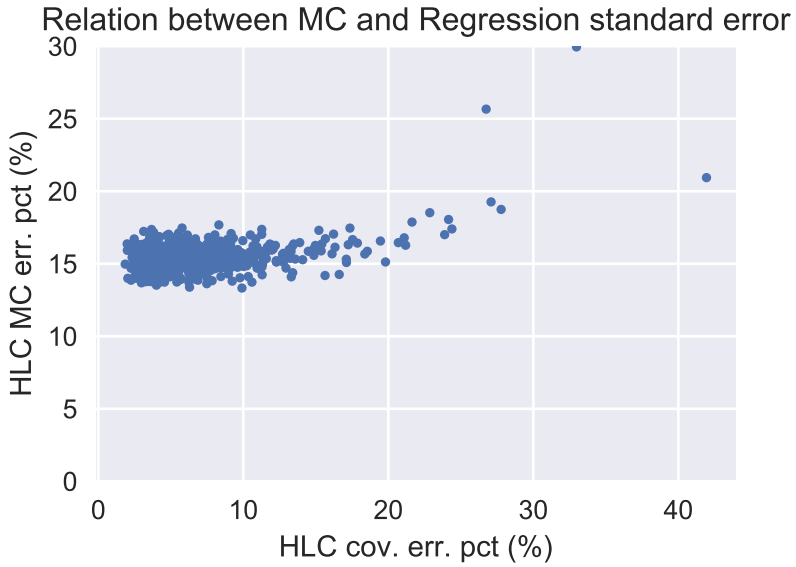

Figure 6.42: Relation between MC derived uncertainty and standard error from regression for 780 EDRP-EDF sites.

is shown in Figure 6.43. The mean standard error for T_0 from the MC approach is 2 ˚ C or 0.7% (of the absolute temperature in Kelvin). This uncertainty is relatively large compared to the typical ranges of interest for internal temperatures.

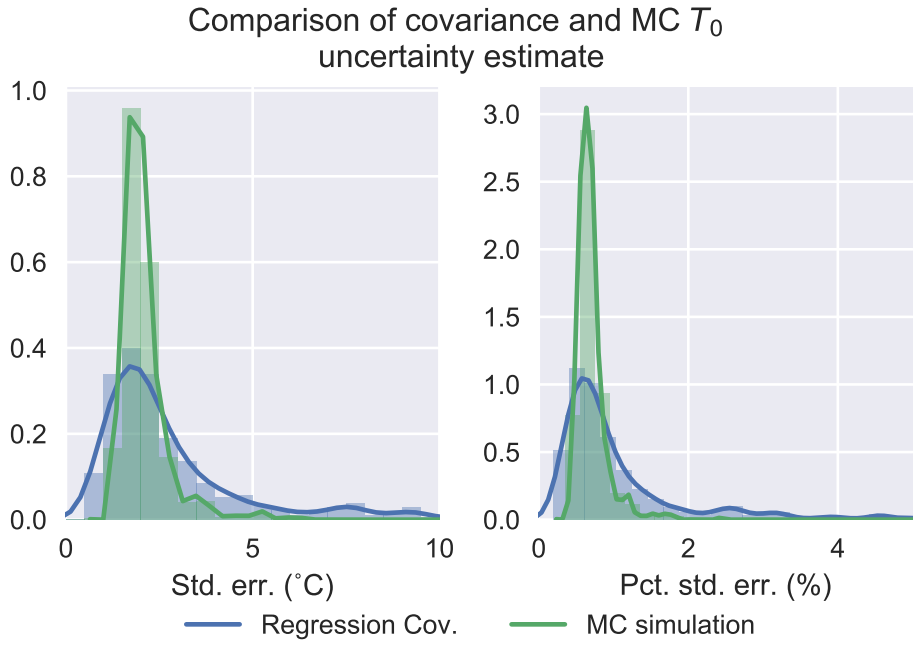

Figure 6.43: Distributions of absolute (left) and percent (of Kelvin value) (right) uncertainty in T_0 for 780 EDRP-EDF sites, using uncertainty derived from covariance and using MC propagation of input uncertainty distributions

6.4 Effect of model assumptions on parameter estimates

6.4.1 Base-load gain

The base-load gain parameter was previously assumed to be 1 ($\eta_B = 1$). This section tests whether this assumption is reasonable, and what impact it has on the T_0 estimates. From eq. 3.27 no effect is expected on HLC, however from eq. 4.4 we can expect the T_0 to depend on the choice of η_B.

Using the EDRP-EDF dataset, the base-load P_B was estimated as described in Chapter 4 Section 4.4.2. η_B was set to a range of values between 0 and 1 and HLC and T_0 estimated in the usual way. The dependence of HLC and T_0 for all sites is shown in Figure 6.44, confirming no effect on HLC and a linear effect on T_0.

Per-site Pearson correlation tests were run between η_B and T_0, and η_B and HLC (Figure 6.45). A linear regression of T_0 against η_B was then performed for each site and the gradients for all sites were found to be equal to $\frac{P_B}{HLC}$. Therefore, as predicted, η_B acts as a linear scaling factor on the internal temperature estimate.

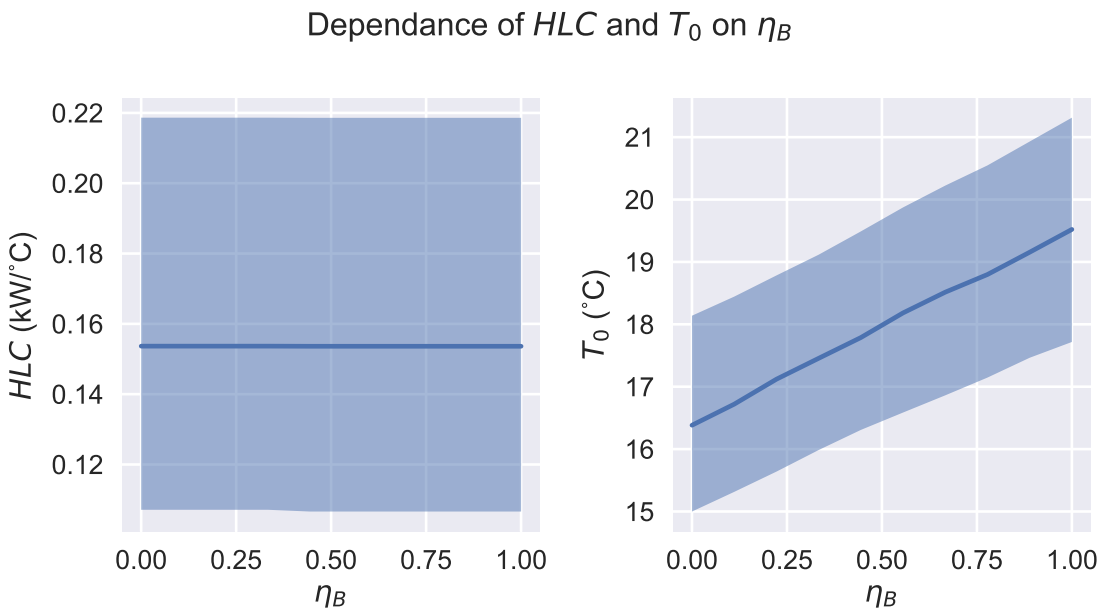

Figure 6.44: Change in estimated HLC and T_0 as a function of η_B across EDRP-EDF. Line shows median value, shaded area 25-75 percentile range

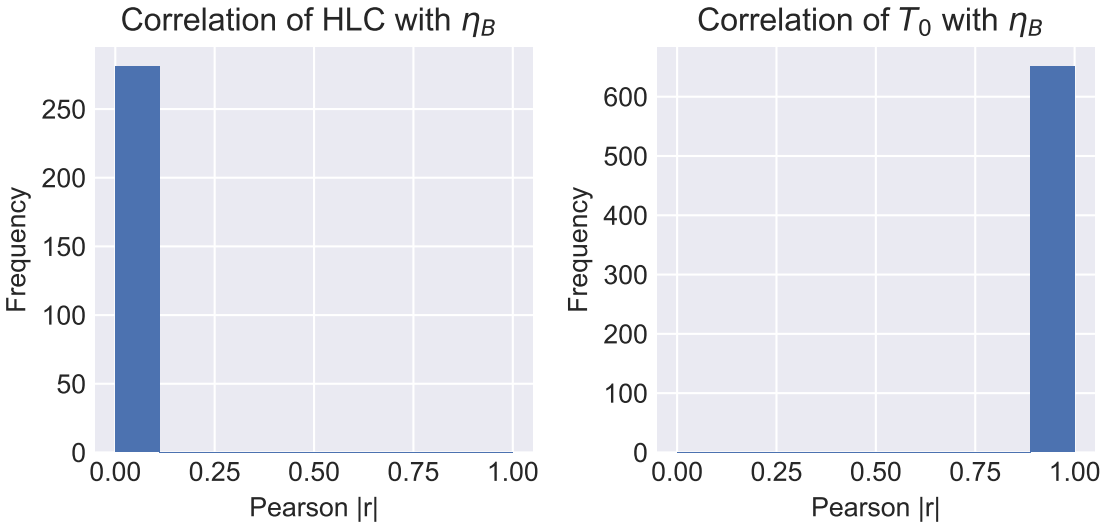

Figure 6.45: Distribution of per site absolute value of Pearson correlation coefficient of HLC and T_0 with eta_B

The amount of change induced by η_B is illustrated by calculating the temperature change between $\eta_B = 0$ and $\eta_B = 1$. The median delta is 3 ˚ C and the distribution is fairly skew (Figure 6.46), this is a relatively large potential change compared to the internal temperature ranges of interest which lie in the 15-25 ˚ C range, and relative to internal-external temperature differences which are also in a similar range. Manual inspection of data revealed that sites most subject to change are those with high scatter, reflecting the greater discrepancy between model assumptions and measurements.

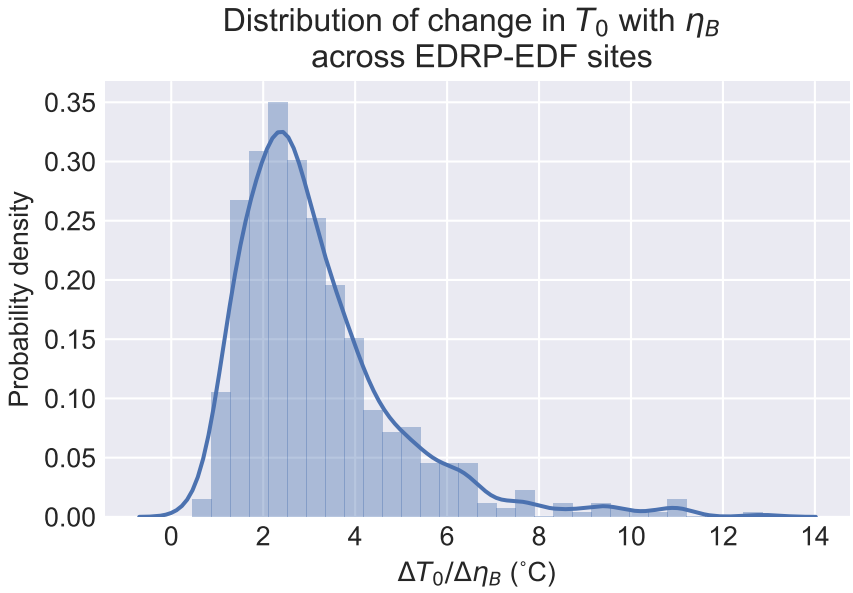

Figure 6.46: Distribution of change in T_0 as a function of η_B across EDRP-EDF.

If η_B can be shown not to vary significantly between dwellings, its effect on T_0 would not be problematic as the ranking of temperatures would be preserved, albeit scaled by an (unknown) η_B. Conversely if η_B varies significantly between dwellings this would result in T_0 not being usable as an indicator of internal temperature.

It is possible to reverse the calculation to estimate η_B given a known T_0. η_B was calculated using the measured internal temperatures from the SWI dataset. Results in Figure 6.47 indicate a problem - many of the estimates lie outside the [0,1] theoretical range for η_B. In effect the parameter is acting as a floating free factor which results in over-fitting. In reality there can be several factors contributing to internal temperature error, as known both from model assumptions and from measurements. Note that using a bounded optimiser limiting the parameter to be in the [0,1] range the distribution suggests that η_B is equally likely to be at the top of the range as at the bottom, which would very much limit the utility of the T_0 parameter.

These results suggest that inferred values of T_0 and η_B through grey-box modelling are unlikely to reflect physical properties of the dwellings. Further work would be useful to estimate η_B from larger power and internal temperature datasets. It may also be possible to measure η_B independently.

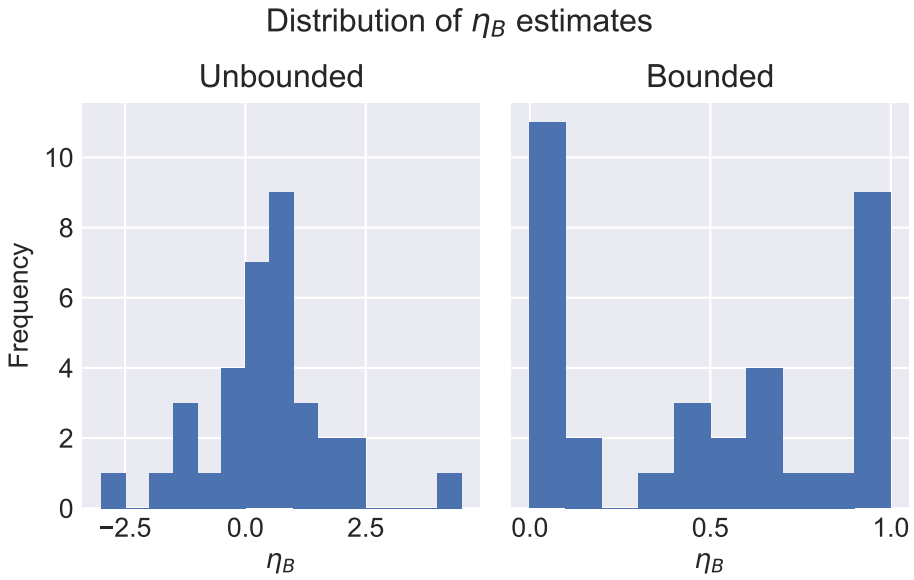

Figure 6.47: Distribution of η_B estimates from measured internal temperatures from SWI dataset using unbounded and bounded solver

6.4.2 Gas boiler efficiency

eq. 3.26 defined the relation between HLC and HTC. HTC provides information about the fabric alone, but can only be calculated if η_{HS} is known. η_{HS} captures both the conversion efficiency of the boiler (i.e. the fraction of the energy content of the fuel converted to heat in water, including Flue Gas Heat Recovery (FGHR)) and the fraction of wasted energy that might still contribute to the dwelling thermal balance (e.g. heat emitted from the body of the boiler unit).

It is interesting to compare HLC and HTC for a dwelling, notably to determine whether relative dwelling performance changes significantly between these two thermal performance indicators. Since the combined effective heating system efficiency is not known, a first order approximation is made by setting the heating system efficiency to the rated boiler efficiency, $\eta_{HS} = \eta_{boiler}$, as defined in the Seasonal Efficiency of a Domestic Boiler in the UK (SEDBUK) database. This assumption ignores

- the known differences between rated and actual boiler efficiency (BRE 2014)
- the difference between boiler efficiency defined as the fraction of fuel energy transmitted to hot water and the effective heating efficiency defined as the fraction of fuel energy transmitted to the building envelope
- the net efficiency when multiple heating systems are used (e.g. gas boiler with additional electric heaters).

The SWI dataset was used as the SEDBUK boiler ratings are included in the SAP data sheets. The SAP main boiler efficiency $\eta_{HS,SEDBUK}$ was applied to the total gas consumption without accounting for cooking use or secondary heating systems. The total thermal power was calculated

as eq. 6.2. To minimise the introduction of errors, the empirical internal temperature model using per site C_T and T_0 derived from internal temperature measurements was used to estimate baseline HLC values.

$$\Phi_H = \eta_{HS,SEDBUK} * P_{gas} + P_{elec} \tag{6.2}$$

The comparison of the SEDBUK boiler efficiency adjusted HLC estimate against the baseline is shown in Figure 6.48. As expected the HTC value is smaller than the HLC value since the metered energy lost through boiler efficiency is accounted for by $\eta_{HS,SEDBUK}$. The efficiency adjusted HLC was on average 0.8 times the baseline (equal to the gradient of the line of best fit on the plot), consistent with the mean SEDBUK efficiency of the sites of 0.8 and confirming the expected result from eq. 3.26.

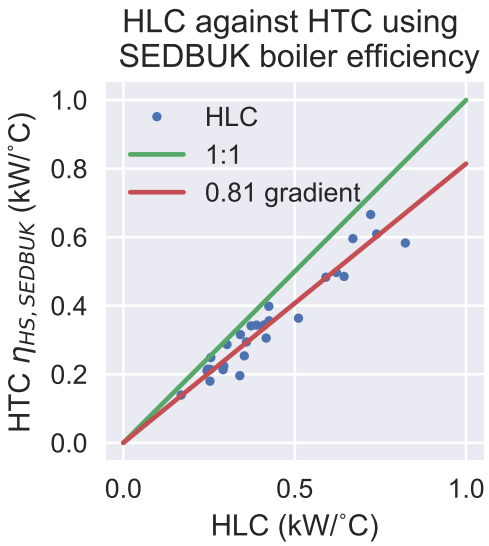

Figure 6.48: Comparison of HTC estimated using SEDBUK boiler efficiencies against HLC estimate, with regression and '1:1' equivalence lines.

Figure 6.49 shows the distribution of differences between HLC and HTC, with a median absolute difference of 0.07 kW/ $^\circ$C and median percentage difference of 17% - although the interpretation of these distributions is made difficult by the relatively small number of sites. For some sites, the heating efficiency has a particularly large impact on the overall performance, leading to differences between HLC and HTC of 30% or more. The possibility of this must therefore be kept in mind when assessing dwellings - some dwellings may be relatively thermally efficient but have very inefficiency heating systems, increasing their overall HLC value.

The potential for dwellings to shift in relative performance is illustrated in Figure 6.50, which shows that while most dwellings have a similar relative performance when measured by HTC or HLC, some dwellings demonstrate large changes in the relative performance ranking.

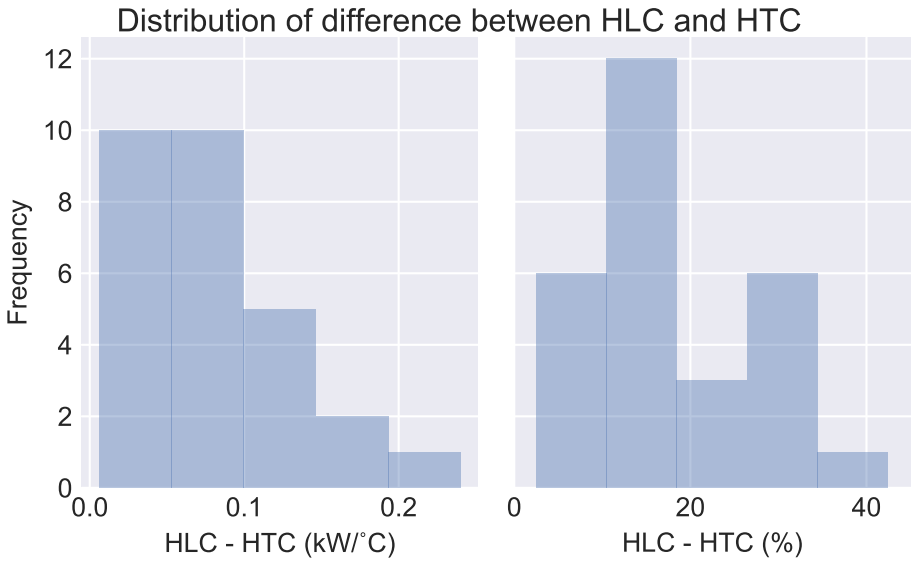

Figure 6.49: Distributions of absolute and percentage difference between HLC and HTC for SWI sites

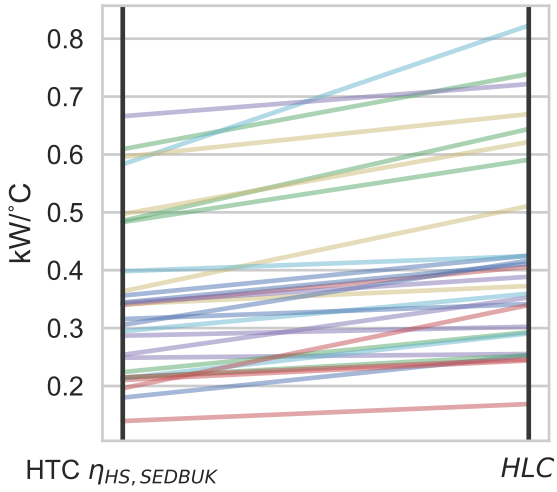

Figure 6.50: Illustration of change in HLC estimate ranking when heating system efficiency η_{HS} is set to boiler efficiency drawn from site survey (SEDBUK rating) compared with assuming $\eta_{HS} = 1$

6.5 Power demand prediction

This section compares HLC estimates made for two consecutive years and the predicted energy demand in the second year using model fit from the first year. The EDRP-EDF dataset has two years of data covering the 2008-2009 and 2009-2010 winters. Year 1 and 2 are defined as:

- Year 1: 2008-08-1 - 2009-08-01
- Year 2: 2009-08-1 - 2010-08-01

The sites are required to have less than 40% missing data and at least 5 valid daily average data points in each year. HLC estimates were made and sites for which no optimal solution was found discarded (these generally reflect corrupted or insufficient data). The results were filtered to only accept sites where CVRMSE $< 30\%$ and the uncertainty (calculated with the Monte Carlo method) was $< 20\%$. The number of sites remaining at each stage is shown in Table 6.7.

Table 6.7: Number of sites accepted by data quality and uncertainty filter.

Filter	Number of sites
Original	780
$< 40\%$ missing, > 5 data points	170
HLC estimate successful	169
CVRMSE $< 30\%$	129
HLC uncertainty $< 30\%$	123

The HLC estimate for year 2 against that for year 1 is plotted in Figure 6.51. There is a good agreement between the two estimates for most sites once sites with excessive CVRMSE and HLC standard error are removed. A single outlier site was removed manually which had a very high HLC in year 1 of 1.25 kW/$^\circ$C, dropping to 0.54kW/$^\circ$C in the second year. Manual inspection of year 1 data for that site revealed it to have unusually high consumption (a large number of days with mean power demand over 20kW), indicating either a fault or a very unusual dwelling situation. With this outlier removed, taking a linear regression of year 2 against year 1 gives a slope of 1.0, indicating a good correlation.

The distribution of absolute and percentage differences between the two years is shown in Figure 6.52. Two sites displayed percentage change magnitude over 100%. It is possible that these sites were subject to retrofits during the measurement period - for example through the installation of improved boiler or boiler controls. Manual inspection of the data indicated that the consumption patterns became more regular in the second year which could support this notion. The median ΔHLC was -0.3% with an inter-quartile range of 27%. After removing these outliers,

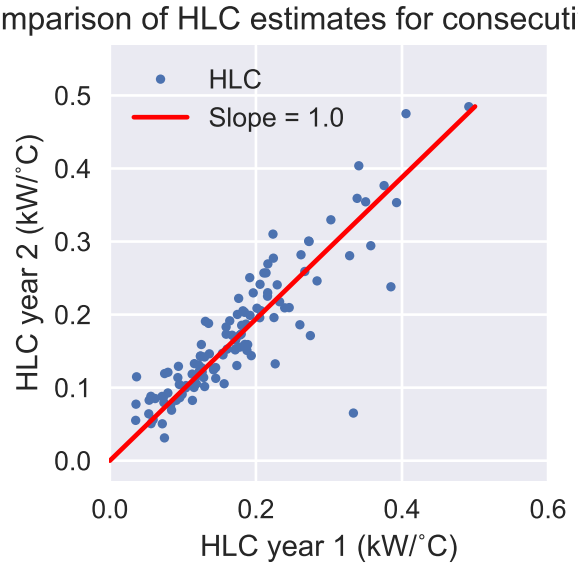

Figure 6.51: Comparison of HLC estimates from year 2 against year 1, with regression line in red

the energy demand in the heating regime (low solar gain period) in year 2 was compared to the value predicted using inferred model parameters in year 1. Figure 6.53 plots the distribution of results and demonstrates a good agreement between them. The distribution of absolute error and percentage error in power are shown in Figure 6.54. A -6% median bias in predicted demand was found with an inter-quartile range of 15%.

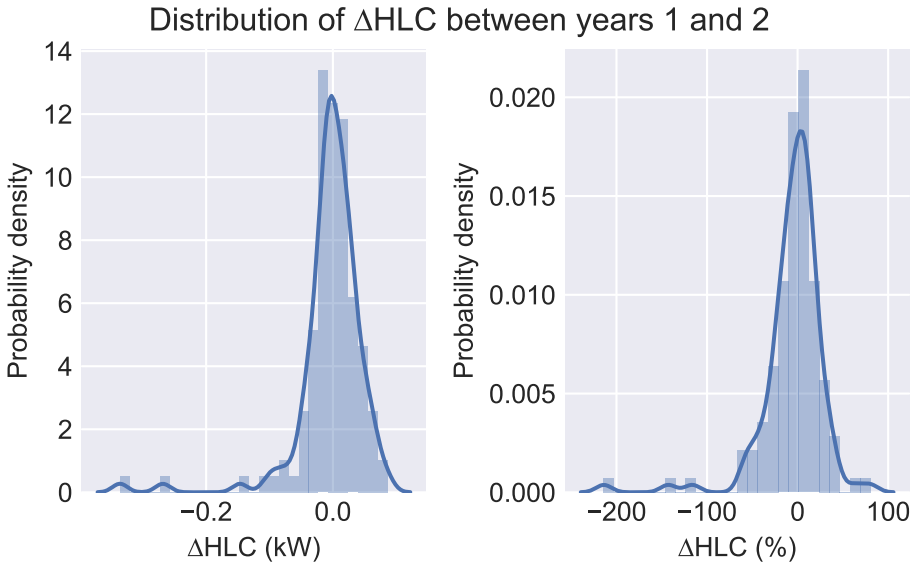

Figure 6.52: Distribution of absolute and percentage difference between HLC estimate for year 2 relative to year 1

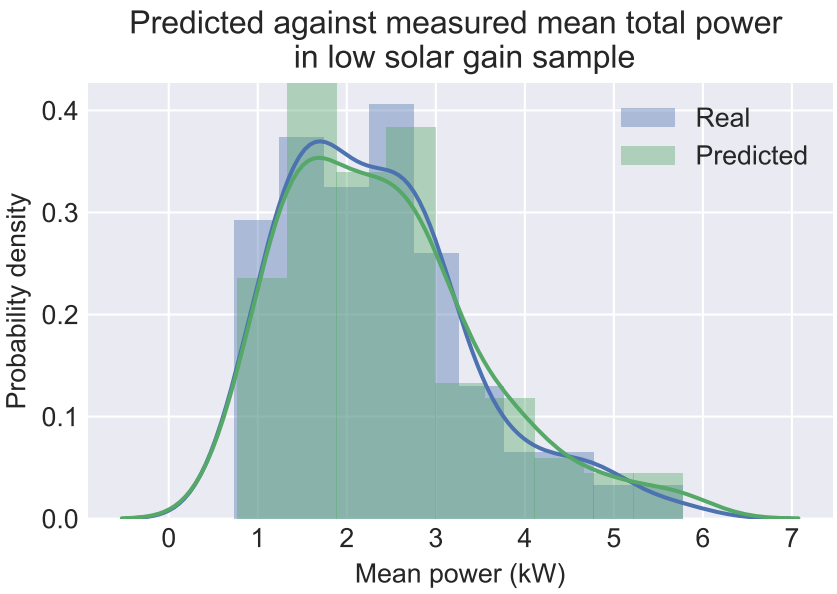

Figure 6.53: Predicted and measured mean power demand in year 2 for the low solar gain sample

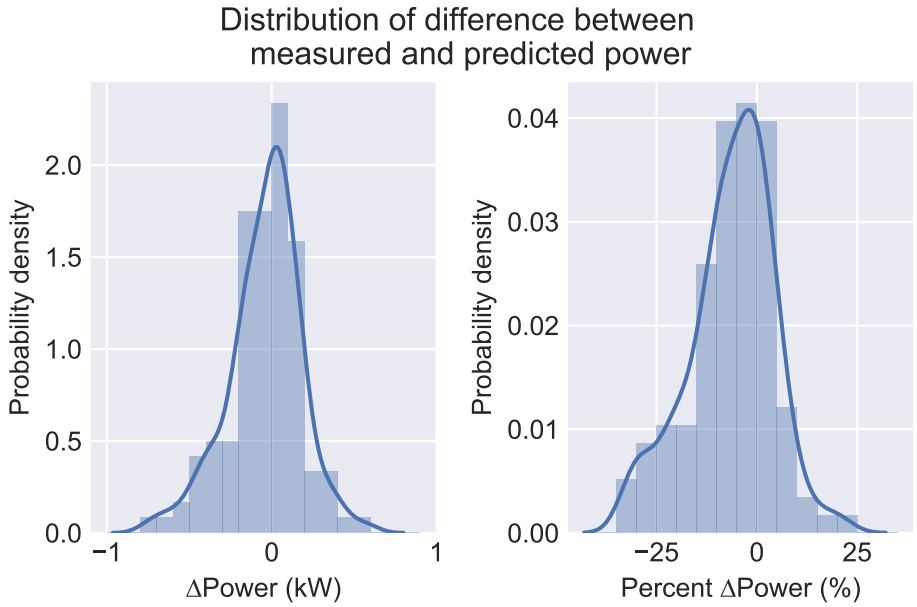

Figure 6.54: Distribution of difference between measured and predicted power, clipped to -1 to 1 range

6.6 Summary

This chapter evaluated a range of aspects of the Deconstruct method. Regarding the data sampling aspects introduced in Section 4.4, it was found that:

- base-load power is independent of external temperature.
- Wind and precipitation do not appear to have a significant impact on total power demand, or on the power-temperature gradient.
- The temperature backstop filter using a temperature cut-off of 15 $^\circ$C is reasonable, having a minimal effect on the power-temperature gradient and reducing the CVRMSE, supporting findings in existing literature.
- The low solar gain filter was demonstrated to reduce the CVRMSE indicating a better model fit, and also avoided a PTG bias.
- The filters also ensure a better model 'pass' rate - i.e. more sites produce an acceptable fit.

Using the EFUS dataset, it was demonstrated that a linear model is a good approximation of the covariance of internal with external temperature. A nationally representative mean of the covariance parameter C_T was found to be 0.17 for England. Dwelling physical parameters - dwelling type and age - were tested as determinants of C_T. It was found there was a small dependence of C_T on dwelling age. Using the SWI dataset, a comparison was made of HLC estimates using mean C_T, dwelling age adjusted C_T, and internal temperature model parameters derived from on-site monitoring. The EFUS mean C_T was found to be a suitable correction, with the only significant improvement coming from fitting the linear model to measured internal temperatures. These results supported using $C_T = 0.17$ for the Deconstruct method.

HLC estimates were found to have an average uncertainty of 15%, in line with existing literature, using the results of a Monte Carlo simulation. Input error distributions were derived from literature. Dwelling and weather datasets were used to derive empirical distributions.

The assumption made in Chapter 4 that $\eta_B = 1$ was evaluated. η_B was found not to affect HLC but to have a potentially large effect on T_0. As a result, T_0 parameter estimate is likely to be an unreliable indicator of dwelling internal temperature.

Using boiler efficiency ratings (η_{HS}), the relation between HLC and HTC was explored using the SWI dataset. It was found that while for most dwellings in the dataset the difference between these two values was under 20%, for a small number the difference could be much larger, up to 40%. Therefore it could be expected that in the majority of cases, indexing dwelling performance by HLC (fabric and heating system losses combined) would produce similar ranking to indexing by fabric performance alone (HTC). Furthermore, the heating system is an important component of the overall dwelling performance, it is therefore desirable to consider when assessing whole dwelling performance.

Finally, the ability of the Deconstruct model to predict energy demand across years was evaluated by using model parameters in one year to predict the demand in the second year. The prediction was found to be good with a median difference in predicted power demand of -0.3%. Applying conservative data quality and output uncertainty filters, it was found that 82% of EDRP-EDF produced an acceptable estimate of HLC, suggesting that Deconstruct should be broadly applicable to existing dwellings. Some of the observed differences in HLC between years might be attributed to physical changes in the dwelling such as retrofits of insulation or heating systems. The predicted mean power demand year-on-year was also found to be within the expected uncertainty range of 15%.

Chapter 7

Results

This chapter presents the results of applying the Deconstruct method to a range of datasets. It is recalled that Deconstruct has been developed as a rapid, scalable method of thermal characterisation for UK dwellings using smart meter data. The chapter is structured as a series of short case studies chosen to assess different aspects of Deconstruct within the constraints of the available data. The results are presented under three sections.

Section 7.1 presents the implementation of Deconstruct at scale. The utility and interest of Deconstruct lies in the ability to be applied to datasets covering a large number of dwellings - essentially all dwellings for which smart meter data is available. A demonstration of this type of application was performed in collaboration with the UK Data Service (UKDS) using the public EDRP dataset. Results were compared to national distributions.

Section 7.2 demonstrates the use of the Deconstruct method to establish a classification of dwellings using Heat Loss Coefficients. Two datasets were used: Pennyland dataset and EDF subset of EDRP. The Pennyland dataset made it possible to investigate the relation between inferred HLC and known construction types, as this project collected data from groups of dwellings purposely constructed to different thermal standards. The EDF dataset presents a more general case comparing the HLC with broad dwellings types. In both cases machine learning unsupervised classification algorithms were applied to separate HLC distributions into groups, and the groups compared with known dwelling properties.

Section 7.3 explores the determinants of the HLC in terms of physical dwellings properties using the EDRP-EDF dataset, and evaluates whether dwelling clusters may be identified from their thermal properties.

7.1 Deconstruct at scale: UKDS Big Data case study

Applying the Deconstruct method to big datasets was one of the aims of the method development and guided its design and implementation. Deconstruct was applied to the 13,868 EDRP sites using a Big Data processing platform developed by the UK Data Service (UKDS), as part of a UKDS project to deliver a proof of concept project using energy data to test and validate the UKDS Big Data platform. For this project, the UKDS provided access to their new cloud computing infrastructure, as well as some project funding. A UKDS member was assigned to provide instruction on the access and use of the system as well as technical support in the translation of relevant software code to the new infrastructure.

The UKDS Big Data analysis platform consists of large compute instances running a PySpark/HDFS software stack. The UKDS intends to develop this into a platform available for researchers to process large new datasets in a variety of fields. Energy data was chosen for a demonstrator project because such data was known to be challenging to work with. The intention was that if the platform could be proven with such data, then other cases would be simple by comparison. The implementation of the data ETL workflow and Deconstruct model Python code was adapted to the cloud computing platform in collaboration with the UKDS Big Data Network Support group. Access to the computing platform helped reduce iteration time on ETL and improved modelling times from days to minutes thanks to the considerable increase in computing power. This was used to apply averaging and outlier filtering to the 8003 sites from EDRP for which both gas and electric data was available.

The weather lookup and averaging workflow described in Chapter 5 Section 5.6 was adapted to run on the service and used to generate daily average weather time series for the Local Authority District (LAD) and NUTS1 regions by which EDRP sites were identified. EDRP sites, which only had NUTS1 region information, required weather to be averaged over a large area (as can be seen in Figure 5.2, most NUTS1 regions are about the size of Wales). This introduced additional uncertainties into the parameter estimates. Unfortunately, it was not possible to obtain more accurate locations for the EDRP dataset.

Performing the weather lookup was particularly challenging due to the high disk read/write demands of the process and the interactions of parallel processes attempting to read and write from a single set of weather files on disk. Given more development time it was envisaged that much more efficient solutions could be devised, however it was decided not to pursue such optimisations for the proof of concept stage. Notably, better support for accessing netCDF/hdf5 files using PySpark/HDFS would be highly beneficial. The weather average was combined with the site power time series and saved into an optimised intermediate tabular file format (Parquet files) suitable for access by PySpark.

Unlike the EDF subset of the EDRP, very little dwelling metadata was available. Most of the

information provided was simple data bookkeeping (start and end dates, number of rows, number of missing points per fuel). The only additional dwelling metadata was an ACORN code. ACORN is a proprietary demographic segmentation system intended primarily for marketing. The code is derived from post codes rather than from surveys of actual dwellings, and as such it does not provide any property specific additional data.

The EDRP data presented a number of challenges, both in terms of scale and quality. Despite the large number of sites in the dataset, of the original 13868 sites only 7529 (54%) passed basic data quality assurance filters, indicating that data collection remains a significant concern for large scale projects. The reduced precision of the site location information increased the fit uncertainties and was a contributing factor to the lower 63% success rate for EDRP dataset sites compared to the 82% success rate for the EDF subset of the same data - the latter had much more precise postcode-level location information.

7.1.1 HLC estimate results

Figure 7.1 and Tables 7.1, 7.2, 7.3 summarise the results of applying Deconstruct to the EDRP dataset. The distribution of the mean values of HLC, P_B and T_0 across the UK NUTS1 regions is shown in Figures 7.2, 7.3, 7.4 and summarised in Table 7.4. These geographic distributions show that despite the size of the input dataset there is still a lack of data for certain regions, caused by the EDRP sampling not being rigorously structured. HLC and base-loads appear to be somewhat lower in the Midlands and Scotland compared to the southern regions, with the South-east showing particular high values, although this may be an artefact of there only being 31 sites in this region.

Table 7.1: EDRP base-load power estimate HLC distribution characteristics

HLC distribution property	Value (kW/ $^\circ$C)
25%	0.086
Median (50%)	0.118
75%	0.160

Table 7.2: EDRP base-load power estimate P_B distribution characteristics

P_B distribution property	Value (kW)
25%	0.293
Median (50%)	0.410
75%	0.558

Table 7.3: EDRP base-load power estimate P_B distribution characteristics

T_0 distribution property	Value ($^\circ$C)
25%	18.2
Median (50%)	20.5
75%	23.7

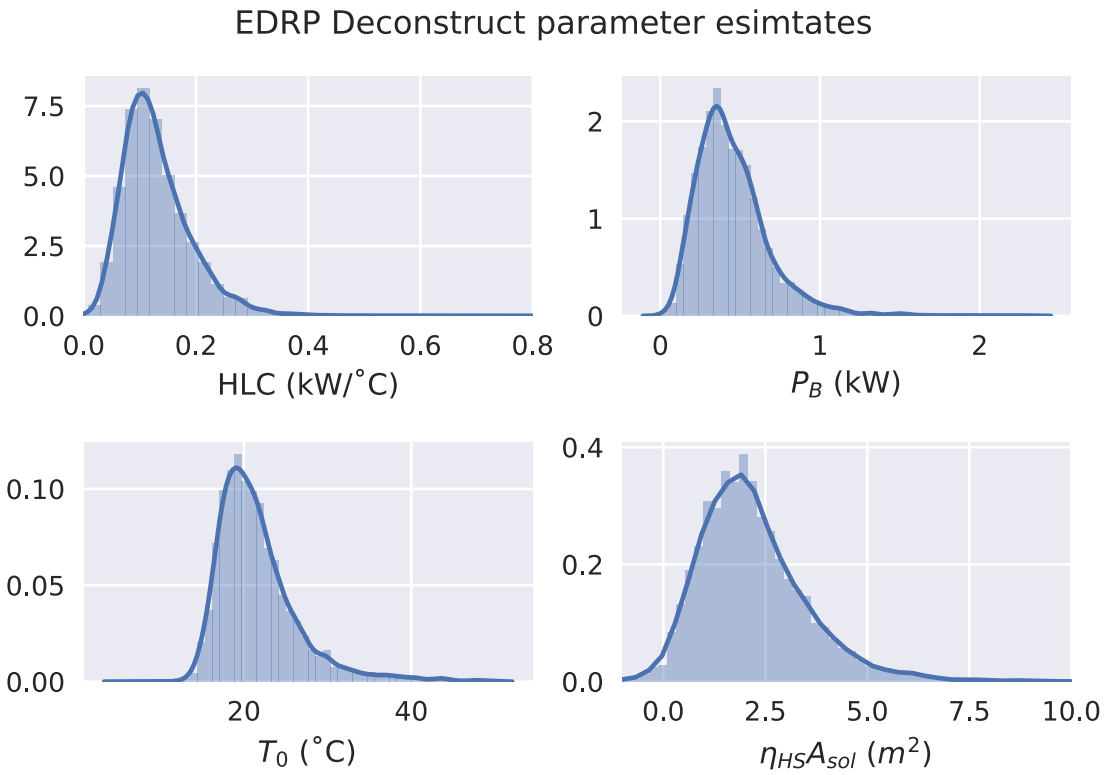

Figure 7.1: Distributions of inferred HLC, base-load power P_B, internal reference temperature T_0, and efficiency-adjusted solar aperture $\eta_{HS}A_{sol}$ for the EDRP dataset

Table 7.4: Summary of HLC, T_0, P_B, and sample size corresponding to Figure 7.2

NUTS region	HLC	T_0	P_B	Sample size
East Midlands	0.12	23.57	0.47	1437
West Midlands	0.12	24.57	0.5	629
Greater London	0.2	20.21	0.54	95
South East	0.16	19.3	0.47	888
South West	0.14	23.77	0.59	31
Wales	0.13	18.98	0.34	310

NUTS region	HLC	T_0	P_B	Sample size
Scotland	0.12	19.91	0.39	1070

With this larger dataset and its somewhat patchwork pedigree, extensive quality control filters were needed. A breakdown of number of sites eliminated at each quality control stage is shown in Table 7.5. Of the 7529 sites which had sufficient input data in the winter months (at least 5 points and a temperature range of 5 $^\circ$C), 4730 produced an HLC estimate with a reasonable level of uncertainty representing a pass rate of 63%. It is important to note that the distribution of HLC estimates is not affected by the removal of high uncertainty sites (Figure 7.5). The HLC 30% uncertainty threshold was based on the covariance uncertainty returned by the optimiser. It was observed in Figure 6.42 using the EDRP-EDF dataset that the Monte-Carlo uncertainty was around 15% for sites where the covariance uncertainty was under 30%. Since calculating an MC uncertainty for the much larger EDRP dataset was too time consuming, the 30% cut-off using the covariance uncertainty was used as a simplification.

Table 7.5: breakdown of site numbers in analysis process

Check	N sites
Total	13868
Has dual fuel type	8523
Has > 0 electric and gas readings	8003
365 days data, $< 50\%$ missing	7038
Winter $T_e x$ range $> 5\,^\circ$C	6979
Optimiser found solution	6952
CVRMSE < 30 %	4530
HLC $\sigma < 30\%$ %	4461

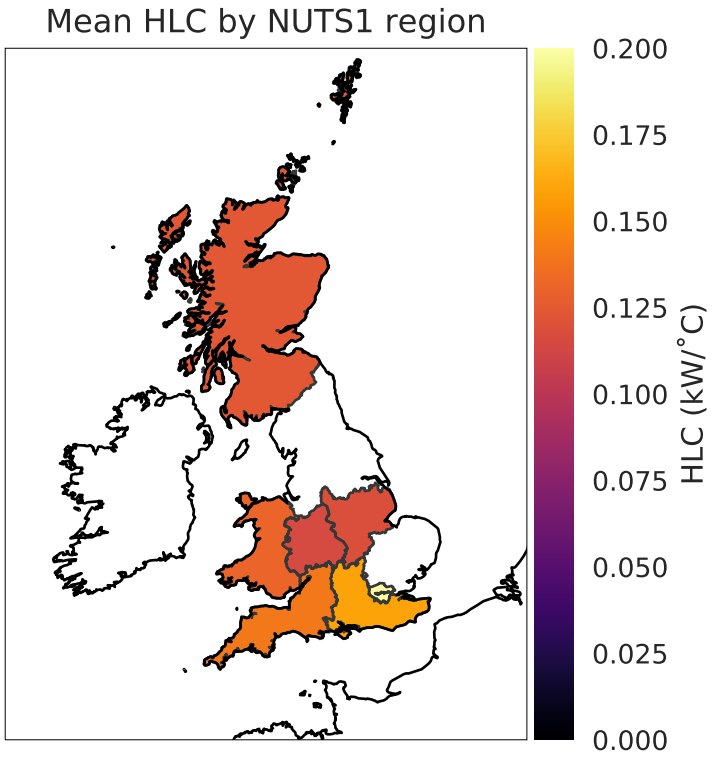

Figure 7.2: Mean HLC by NUTS1 region

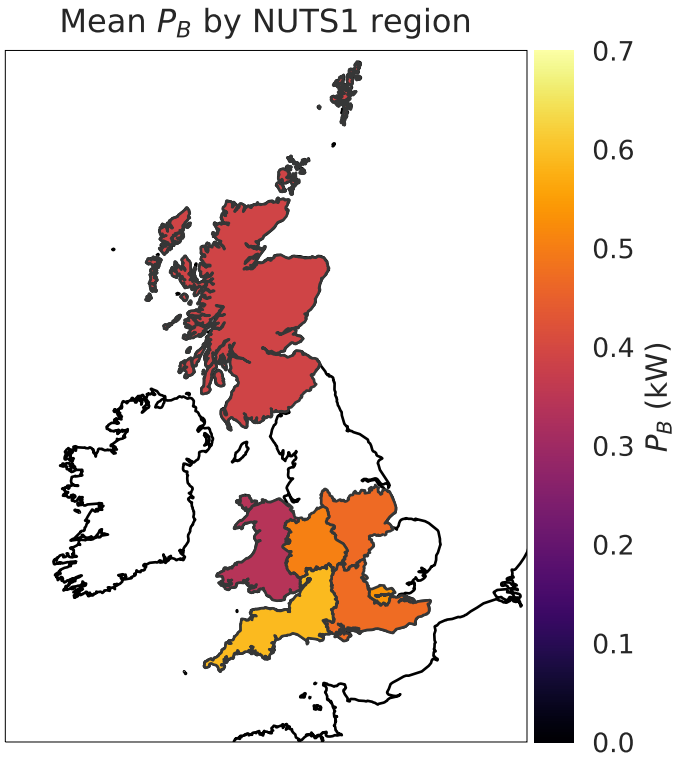

Figure 7.3: Mean P_B by NUTS1 region

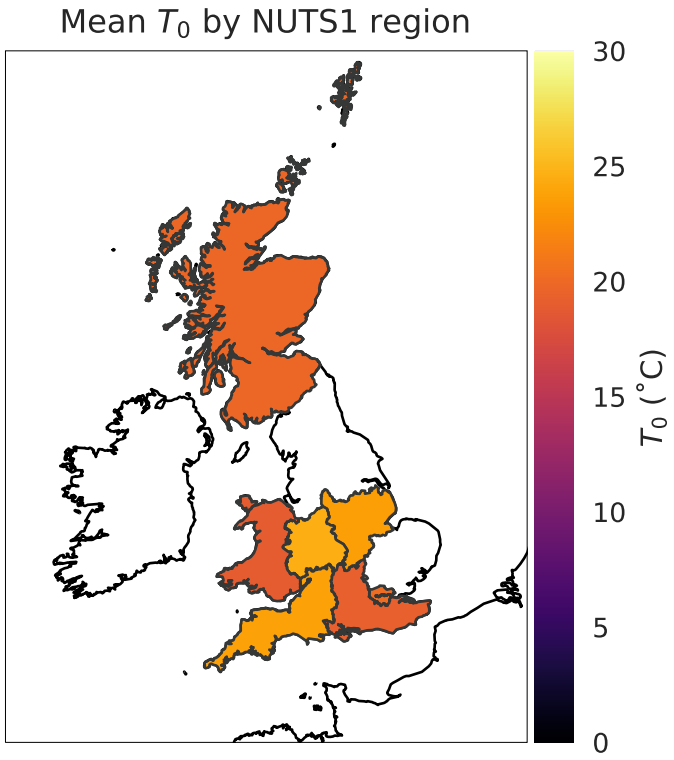

Figure 7.4: Mean T_0 by NUTS1 region

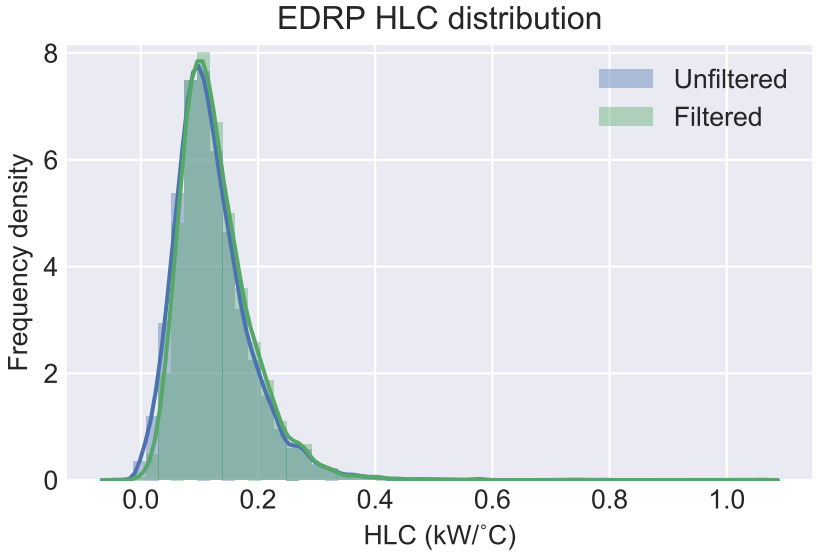

Figure 7.5: Distributions of HLC estimates before and after removal of high CVRMSE and high uncertainty sites

The pass rate is somewhat lower than that observed for the EDF subset. This may be due to two possible factors. Firstly, the full EDRP contains a greater diversity of dwellings and therefore includes more sites for which the simple thermal model is not a good fit. This is reflected in the 2572 sites eliminated by the CRMSE cut-off, which essentially determines whether the model can be considered a reasonable representation of the real energy demand. In the future, it would be interesting to more closely study the dwellings which fail this test. It may be that an alternative thermal model would more closely represent their energy demand pattern. The second factor to consider is that the EDRP had less accurate location information (LA codes and NUTS1 region codes) compared to the EDF subset (postcode district or areas). This would have a corresponding impact on the accuracy of the weather variables as they must be averaged over a larger area.

7.1.2 Comparison of data-driven and bottom-up approaches

The Deconstruct method could form the basis of an 'empirical' EPC, defined by metered energy demand instead of bottom-up modelling. It is therefore interesting to compare the Deconstruct results against existing (Rd)SAP based approaches by comparing normalised yearly energy demand estimates. This was performed by comparing a Normalised Annual Consumption (NAC) derived from the Deconstruct results for the EDRP dataset against the UK EPC database and the Cambridge Housing Model (CHM).

The database of UK EPC ratings was released in April 2017 (DCLG 2017). It includes EPC ratings and associated metadata for approximately 16 million dwellings and includes a current energy consumption value. There are however a number of caveats regarding this value:

- The EPCs do not identify which revision of RdSAP was used to generate the rating - as

RdSAP has evolved through the years there can be some differences in the method and resulting energy rating.

- Energy consumption is primary energy use. This mainly impacts the electric fraction of total energy demand, as the final electric energy is converted to primary fuel use based on standard efficiencies and energy generation fractions for the UK energy mix. This is therefore somewhat different to total metered demand and somewhat subject to changes in the energy mix coefficients across the different RdSAP revisions present in the EPC database.

Figure 7.6 shows the distribution of domestic EPC ratings for 2010-2012 (calculated as the sum of ratings for each band across those years) (Department for Communities and Local Government 2017).

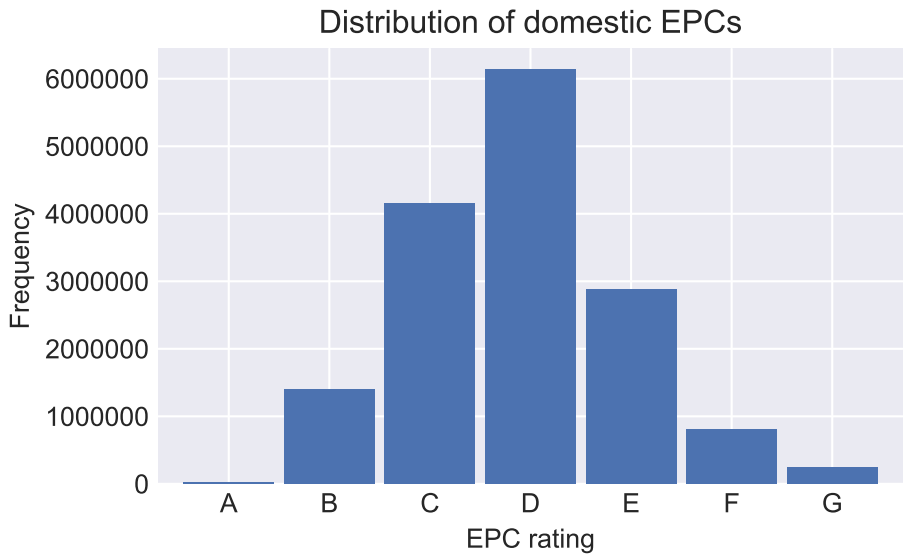

Figure 7.6: Distribution of the domestic EPC ratings for 2010-2012

The Cambridge Housing Model was developed by Cambridge Architectural Research for DECC to underpin the 2012 "Housing Energy Fact File" and "Energy Consumption in the UK" reports. It uses EHS 2011 data, coupled to a SAP-based energy calculator, to estimate energy use for all homes in England, broken down by final use (Palmer, Tillson, and Armitage 2013; Davey 2016). The Excel-based tool generates yearly energy demand values for permutations of SAP input parameters weighted by the estimated prevalence of each permutation. The resulting weighted distribution aims to be representative of national yearly energy demand.

7.1.2.1 Normalised Annual Consumption

In order to compare the EDRP results with nationally representative distributions, it was necessary to calculate a Normalised Annual Consumption (NAC) (kWh). However, in order to calculate NACs it was necessary to calculate a solar gain parameter in order to calculate mean power for every day of the year according to eq. 3.20 and eq. 4.1. The solar gain parameter was calculated

as $\frac{d\Phi_{sol}}{dI_{sol}} = \eta_{HS} A_{sol}$, where $\Phi_{sol} = P_{tot} - HLC(T_{in} - T_{ex}) - P_B(1 - \eta_B/\eta_{HS})$ using values of HLC and P_B estimated as described previously. Note that the choice of assumption for η_B and η_{HS} is not important for the solar gradient. A daily data sample was selected applying:

1. An error filter (Section 4.4.1.1).
2. A temperature cut-off filter eliminating days where T_{ex} is greater than 15 °C (Section 6.1.4).
3. Outlier removal filter (Section 4.4.1.2).
4. Unoccupied days filter (Section 4.4.1.3).

The $\eta_{HS} A_{sol}$ coefficient was estimated from the resulting sample.

To calculate NAC, the daily power was calculated by plugging the resulting model parameters into eq. 3.20 and eq. 4.1 using a long term weather average. The weather average was selected as the 20-year mean daily weather for the North East region from 1987-2008, using CFSR data. This period was selected to align with the SAP2009 weather averages which correspond to the 20-year average for the East Pennines region, 1987-2008 (Wingfield 2017). In the absence of an exact definition of 'East Pennines' within the SAP method, the North East Government Office Region was used. Annual total power demand in kWh was calculated from daily demand.

Note that the global horizontal solar irradiance (GHI) was used for the $\eta_{HS} A_{sol}$ estimate. Stamp (2015) notes that different solar projections, such as on a south-facing vertical plane, may be more appropriate for accurately calculating solar gains. GHI was retained as it corresponds to the values reported by the CFSR weather models and should be suitable for a first order approximation.

Finally, due to the lack of suitable dwelling metadata in the EDRP dataset, it was not possible to compare NAC against EPCs for the same dwellings. Had the dwelling address data for EDRP not been censored for privacy reasons, the dwellings could have been matched with those in the EPC database using the address data. It was therefore only possible to compare EDRP NAC distributions against national EPC distributions.

7.1.2.2 Yearly energy demand comparison

The comparison between the yearly energy demand distributions for the EDRP measured yearly demand, EDRP NAC, CHM model, and EPC archive is shown in Figure 7.7. The measured and NAC distributions show considerably smaller demands than the CHM and EPC results based on bottom-up, RdSAP driven energy models. This suggests that the latter models may be systematically overestimating dwelling demand, although there are a number of important caveats. As noted previously, the EPC distribution is for primary energy and is therefore expected to have a larger median. It is therefore perhaps more informative to focus on the CHM distribution.

Furthermore, the EDRP dataset is not necessarily nationally representative and unfortunately lacks the requisite metadata to either check how representative it is or weight the resulting distri-

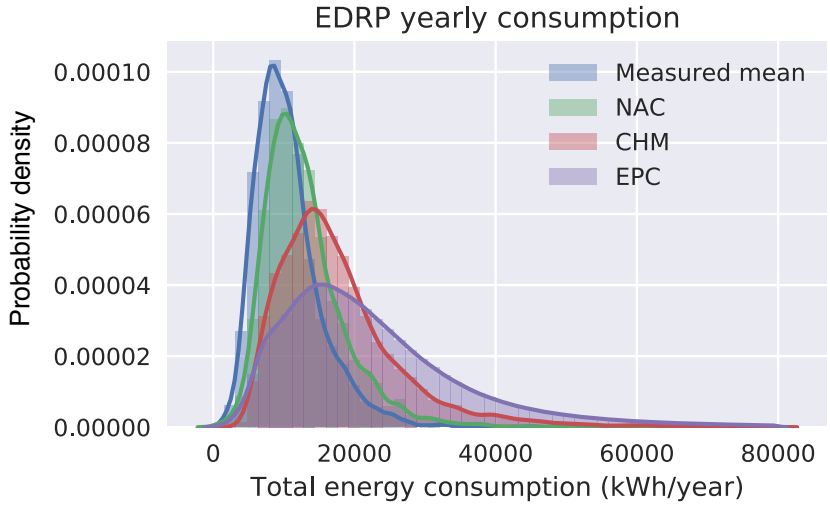

Figure 7.7: Distribution of mean yearly power demand for EDRP sites compared to the NAC calculated using Deconstruct parameters, the CHM model output distribution, and the public EPC database total energy consumption distribution

bution accordingly. It was originally surmised that the relatively large sample size in the EDRP dataset would result in it being relatively representative of UK dwellings. However, calculating NAC for the EDRP-EDF subset, which has more accurate location data (and therefore weather) as well as showing a distribution of house age and type similar to the UK mean (Chapter 5, Figures 5.12, 5.13) shows a much closer agreement between NAC and CHM (Figure 7.8).

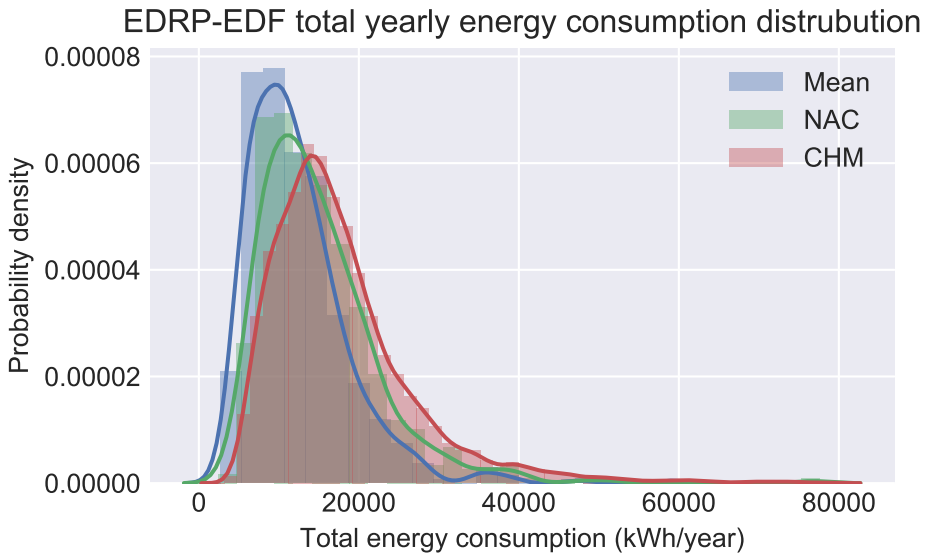

Figure 7.8: Distribution of mean yearly power demand for EDRP-EDF sites compared to the Normalised Annual Consumption calculated using Deconstruct parameters, and the CHM model output distribution

Given the known limitations of bottom-up SAP-based energy demand estimates, it is not possible to determine which of NAC or CHM values are more accurate with the available data - although clearly the NAC more closely represents the mean measured consumption from which it was derived.

There has been some discussion (BEIS 2016c) on how empirical performance metrics, such as HLC, could be converted into an 'empirical EPC' rating. These results indicate that at the very least, HLC or the derived NAC produce figures within the expected range. It also demonstrates that there is no 'natural' grouping of performance into bands as is done in the EPC 'A' through 'G' as the distributions from different dwelling efficiency types overlap. This will be further explored in Sections 7.2, 7.3.

7.1.3 Comparison of T_0 with EFUS distribution

Using the EFUS internal temperature dataset the expected national distribution of internal reference temperatures (T_0) can be derived by fitting the simple linear temperature model to all sites and comparing the weighted distribution of these to the distribution of T_0 estimated from EDRP. The distributions are compared in Figure 7.9. T_0 was previously found to be sensitive to the choice of base-load gain parameter η_B, these results are calculated assuming $\eta_B = 1$. The effect of the base-load gain on the internal temperature distribution is very clear in Figure 7.10. Note that as established previously, the choice of η_B has no effect on the HLC or P_B parameter.

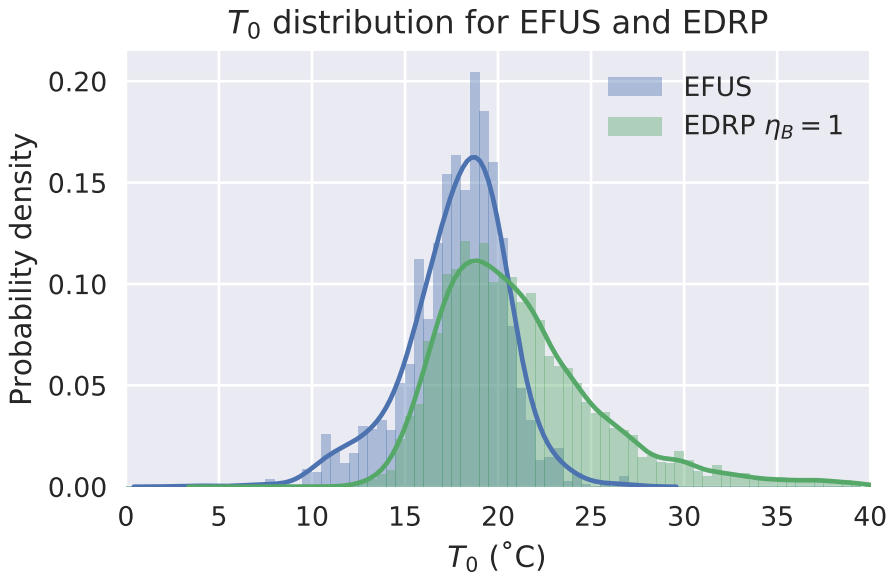

Figure 7.9: Distribution of T_0 estimated from EFUS internal temperatures compared to T_0 distribution inferred from EDRP sites

The inferred T_0 distribution is not the same distribution as that found by EFUS (confirmed by a K-S 2 sample test of distribution similarity). It can be seen that that the inferred distribution is right-skewed while the empirical one is somewhat left-skewed - that is, the inference method tends to generate a 'fat tail' of high internal temperatures compared observed temperatures. Figure 7.10 illustrates that the distribution is shifted by the choice of η_B, without fundamentally changing its shape. It suggests the possibility of using nationally representative distributions to calibrate the Deconstruct model assumptions, by finding the η_B value which causes the best match between

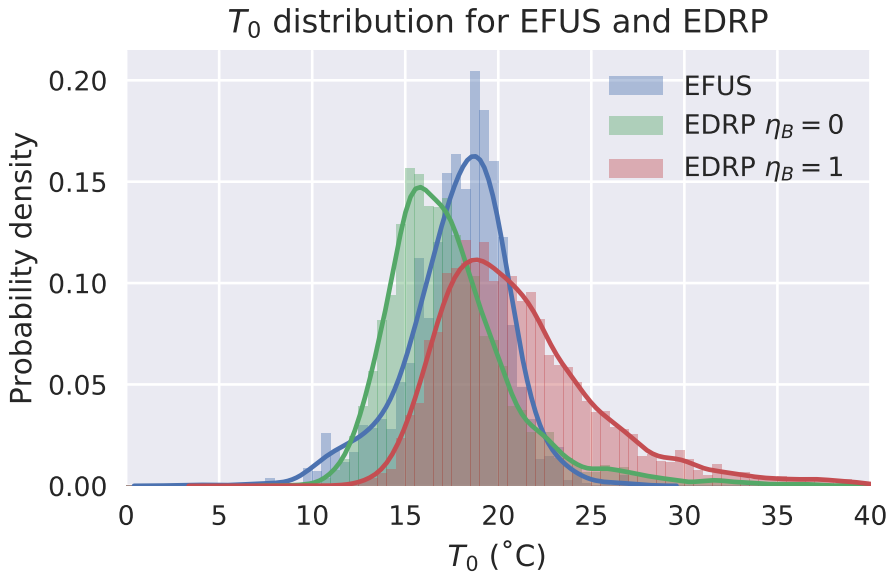

Figure 7.10: Distribution of T_0 estimated from EFUS internal temperatures compared to T_0 distribution inferred from EDRP sites showing distributions for η_B limit values 0 and 1.

distributions. This result also confirms that the inferred T_0 value a relatively unreliable indicator of measured internal temperature. This is further discussed in Section 8.1.3.

7.2 Classification of dwellings using HLC: Pennyland case study

7.2.1 Pennyland sites thermal properties

The Pennyland case study (see Chapter 5 Section 5.5.2) investigated the effect of different dwelling design and construction approaches by constructing different groups of dwellings on the estates to different standards and comparing the resulting energy demand. Three variations were built: one to contemporary 1975 British building standards, one to improved insulation standards (the at the time prospective 1982 standard) and one to further improved standards with the addition of passive solar design principals (roughly corresponding to Danish 1979 standards) (Table 7.6). The uniformity of build types in each group eliminates an important source of variability, making it possible to assess the relation between power demand and thermal performance, as well as between variants of the method of estimating thermal performance. It should be noted that each building group was equipped with a single boiler type per group thereby reducing the expected impact of the variability in boiler efficiency within each group.

Table 7.6: Number of buildings in dataset for each construction type

Area	Insulation level	N sites	Site Ids
Neath Hill	Baseline	18	1-18
Pennyland Area 1	High insulation	39	19-110
Pennyland Area 2	Passive solar	19	112-130

Figure 7.11 shows the power demand distribution for all sites, while Figure 7.12 splits the distribution by the known building group. The low insulation Neath-Hill is strongly separated from the high efficiency dwellings, except for a small number of outliers. The two high-efficiency groups display significant overlap, making it difficult to distinguish the groups by measured power demand alone. This was already noted in the original report (Chapman, Lowe, and Everett 1985). Nonetheless the distributions do have statistically significant different medians (Kruskal-Wallis test, p-value $<< 0.05$).

HLC estimates were made using the Deconstruct method using measured internal temperatures and with the linear internal temperature model (C_T from EFUS). Estimates were made with both methods as the dataset contains only weekly measurements, losing a lot of information relative to daily readings. Figure 7.13 clearly shows that the modelled variant has a bias in the HLC estimate larger than that displayed in the similar analysis on the SWI dataset in Chapter 6. This is further reflected in Figure 7.14 and Figure 7.15, which display a much clearer grouping of HLC values when local temperatures are used.

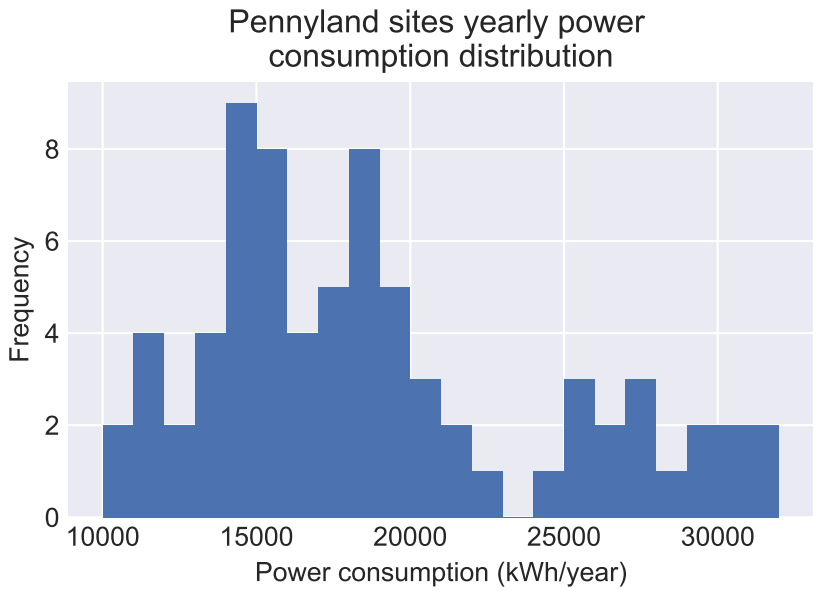

Figure 7.11: Distribution of yearly power demand for sites in the Pennyland dataset

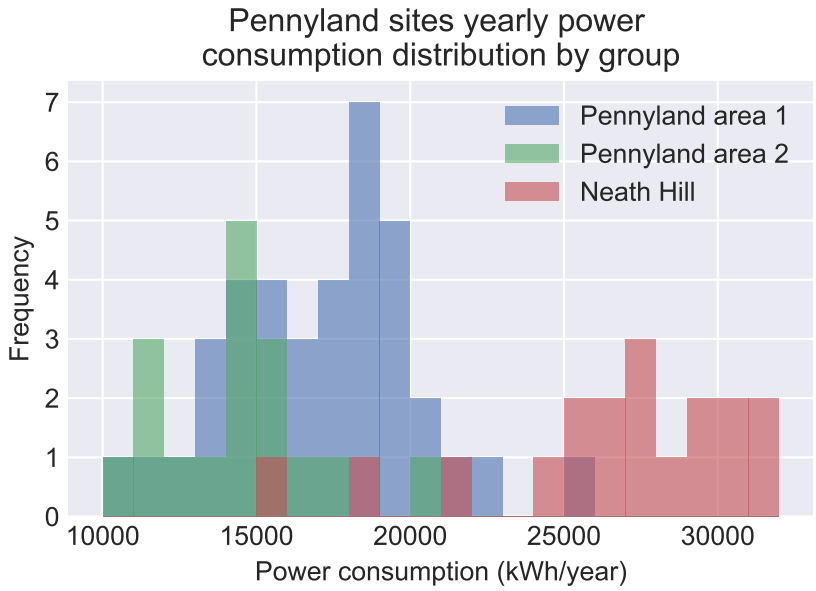

Figure 7.12: Overlaid distributions of mean yearly energy demand for Pennyland areas 1 and 2 and Neath Hill

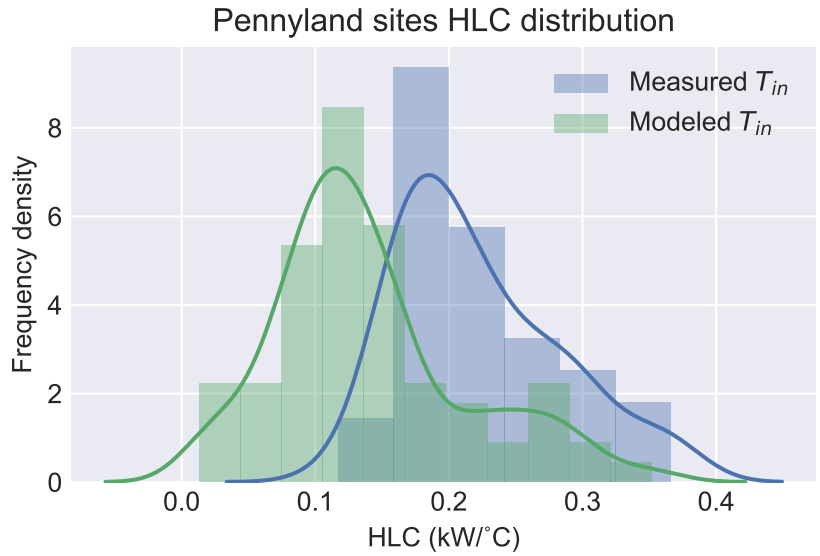

Figure 7.13: Distribution of HLC estimates using measured internal temperatures

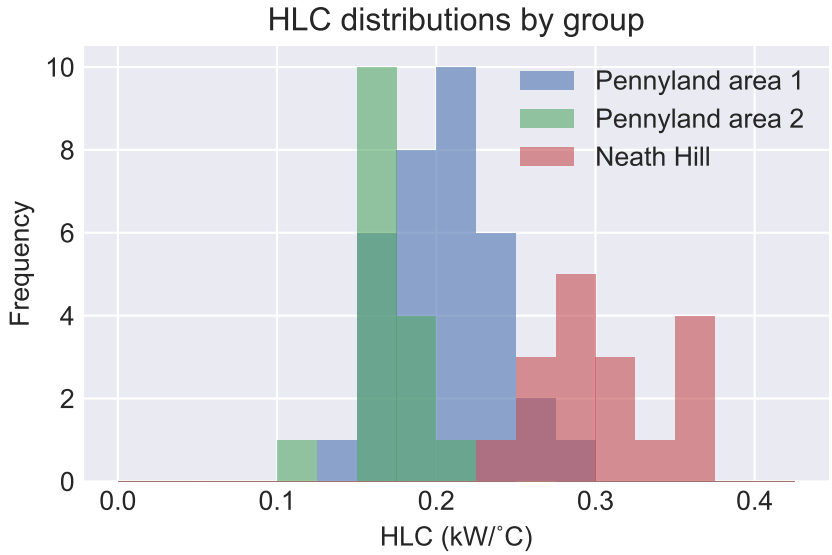

Figure 7.14: Distributions of HLC estimates using measured ΔT for the Pennyland areas 1 and 2 and the Neath Hill control sites

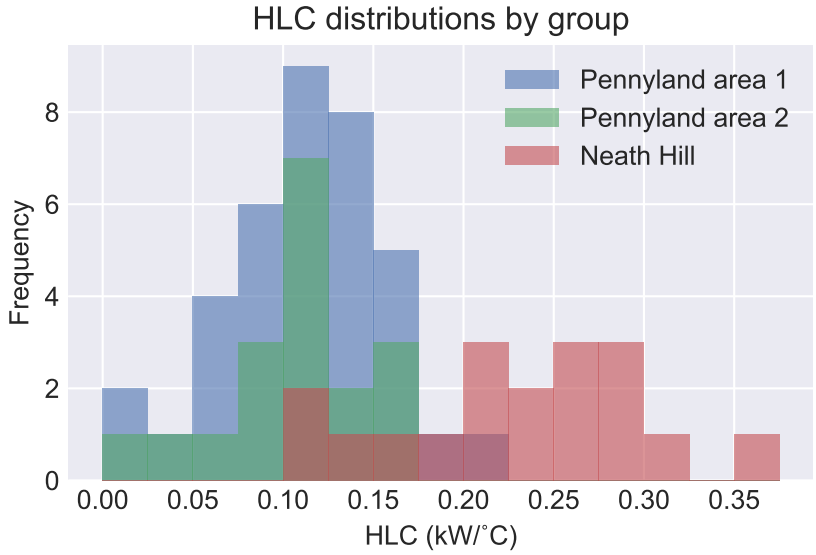

Figure 7.15: Distributions of HLC estimates using modelled internal temperatures for the Pennyland areas 1 and 2 and the Neath Hill control sites

7.2.2 Dwelling classification

Machine learning methods may classify sites into performance bands and enable comparison of the classification result with known dwelling types. For this type of problem of separating overlapping distributions, a Gaussian Mixture Model (GMM) is suitable (EPFL 2013). A gaussian mixture model $P(x)$ is defined by a sum of gaussians with means μ and covariance matrices Σ:

$$P(x) = \sum_i w_i \, \mathcal{G}(\mu_i, \Sigma_i)$$

The GMM classifier uses the expectation-maximisation algorithm to estimate the mean and standard deviation of a predetermined number of distributions mixed in a dataset (Dempster, Laird, and Rubin 1977). This implies that the number of expected groups must be chosen first, either through other knowledge of the dataset or using a separate algorithm. The "GaussianMixture" model class from the Scikit-learn toolkit was used as a ready implementation of this approach (Pedregosa et al. 2011).

7.2.2.1 2-group classification

In the first instance, only two efficiency classes were considered since both Pennyland areas 1 and 2 were built to similar standards, with Area 2 focusing on solar gains improvements which would not affect the results when using the low solar gain sampling. The GMM classifier was used to estimate the distribution components of each class and assign each dwelling to its appropriate class. The class with the lower mean HLC was annotated as the high efficiency dwelling group

and the and high HLC annotated as the low efficiency dwelling group. This was compared with the known groups, with sites 1-18 being the Neath Hill low efficiency group and sites 19-130 being the Pennyland high efficiency group.

The GMM classifier does not ascribe meaning to the generated groups (i.e. it does not 'know' that a given group ID should correspond to the high or low efficiency buildings). The groups were identified as corresponding to the high or low efficiency class based on the estimated group means (high power or heat loss corresponds to low efficiency, and vice versa).

The summary of the HLC and mean power demand for the low and high efficiency groups can be found in Table 7.7.

Table 7.7: Base properties of the Pennyland dataset

	N sites	HLC (kW/ $^\circ$C)	Mean power (kW)
Low efficiency group	17	0.31	3.01
High efficiency group	50	0.19	1.82

As a baseline, the dwellings were classified using the dwelling mean power distributions, which was shown in Figure 7.12 to be strongly associated with the dwelling group. The classification results as shown in Table 7.8 deliver a good agreement between the expected and actual derived efficiency classes (97% of sites correctly classified, 0 incorrectly classified as low efficiency, 2 incorrectly classified as high efficiency).

The HLC estimates were then used to classify dwelling performance. As can be seen from Figure 7.14, there is a less clear distinction between the HLC values compared to the mean power. The results presented in Table 7.9 indicate a slightly reduced 94% correct classification rate, with 4 sites incorrectly classified as low efficiency and none incorrectly classified as high efficiency. This result suggests that HLC corresponds to the dwelling performance class and therefore may enable identification of dwelling performance class from metered data.

Table 7.8: Results of GMM classifier on mean power

	N sites	HLC (kW/ $^\circ$C)	ΔHLC (%)	Mean power (kW)	ΔMean power (%)
Low efficiency group	15	0.31	-0.2	3.19	3
High efficiency group	52	0.20	-2	1.85	-0.5

Table 7.9: Results of GMM classifier on HLC calculated using measured internal temperatures

	N sites	HLC (kW/ ° C)	ΔHLC (%)	Mean power (kW)	ΔMean power (%)
Low efficiency group	21	0.30	3	2.93	5
High efficiency group	46	0.19	3	1.80	2

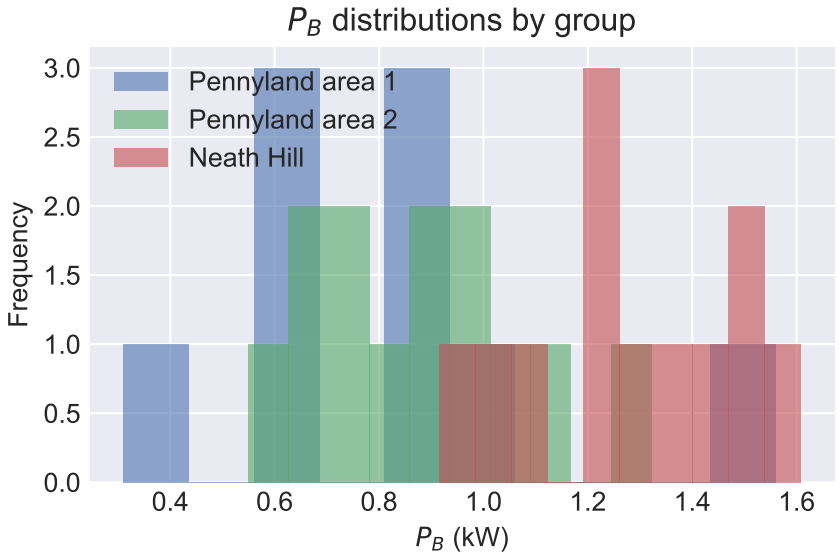

Figure 7.16: Distributions of P_B estimates sing measured ΔT for the Pennyland areas 1 and 2 and the Neath Hill control sites

It is important to note that the overall performance difference between the high and low efficiency dwellings was more marked when using mean power than when looking at HLC. It is possible that differences in boiler technology between groups account for this difference, as the high efficiency dwellings were all equipped with low thermal capacity heating systems while the low efficiency ones all used traditional gas boilers with water storage tanks - this is likely one cause of the significant difference in base-load power, illustrated in Figure 7.16. The high base-load results in higher power demand overall for the Neath Hill sites.

The classification was repeated using the HLC estimated using modelled internal temperatures (i.e. using external temperature data only), with results presented in Table 7.10. This approach correctly classified 92% of sites, with 1 site incorrectly identified as low efficiency and 5 sites incorrectly identified as high efficiency. The classification performance therefore remained good despite the limitations in terms of using weekly data and modelled internal temperatures. The mean HLC of the low efficiency group was underestimated by 13% compared to the true group mean, which was larger than with previous approaches.

Table 7.10: Results of GMM classifier on HLC using modelled internal temperatures

	N sites	HLC (kW/ ° C)	ΔHLC (%)	Mean power (kW)	ΔMean power (%)
Low efficiency group	14	0.26	-13	3.10	-3
High efficiency group	59	0.11	-0.6	1.90	-3

7.2.2.2 3-group classification

The GMM classifier was run on 3 groups using measured internal temperatures, corresponding to the two Pennyland high efficiency variants and the Neath Hill low efficiency sites. Due to the overlap of the distributions, the results are less stable (i.e. the classifier output sometimes produces different results). Notably, while the group means are relatively stable across runs, assigning a given dwelling to a particular group is less stable, resulting in the group assignment prediction being reduced and more variable across runs. Table 7.11 presents a typical run result. As expected from the distributions, the classifier struggled to distinguish between the similar Pennyland Area 1 and 2 sites, although the group HLC mean values were close to their true value (the greatest difference being 7% for Pennyland Area 1). These results suggest that identification of dwelling groups using the classification approach is likely to be less successful when groups are not very clearly defined.

Table 7.11: Mean HLC estimated from internal-external temperature difference for each group

Area	N sites	GMM N	HLC (kW/ ° C)	GMM HLC (kW/ ° C)	ΔHLC (%)
Neath hill	17	19	0.31	0.30	0.8
Pennyland Area 1	34	21	0.21	0.22	-7
Pennyland Area 2	16	27	0.17	0.17	1

7.3 Determinants of HLC estimates: EDRP-EDF case study

7.3.1 Dwelling HLC by type and age

It is interesting to consider the determinants of HLC estimates in terms of the physical properties of dwellings such as age and dwelling type. Work along these lines was performed previously in Summerfield, Oreszczyn, et al. (2015a) using PTG estimates on the EDRP-EDF dataset (which is convenient for the presence of suitable dwelling metadata).

The Deconstruct method was used to estimate HLC for EDRP-EDF data and the results associated with site metadata, which included dwelling type (detached, flat, etc) and dwelling age band, as well as information such as number of bedrooms. This enables a comparison of dwelling HLC as a function of basic physical properties - type, age, and size.

536 dwellings were retained as having suitable input data with CVRMSE $< 30\%$ and HLC standard error $< 30\%$. The yearly mean power demand for those dwellings was also calculated in order to compare power and HLC distributions - the dwelling thermal loss is known to be a significant driver of total demand so it is useful to verify whether these distributions follow the same trends.

Figure 7.17 shows the power and HLC grouped by dwelling type. Different dwelling types (detached, semi-detached, flats, etc.) have different thermal property distributions - for example flats tend to be more efficient than detached houses due to differences in ratio of exposed surface area to dwelling volume. HLC is smallest (and has the smallest range) for flats, while detached houses have both the largest median and the largest minimum HLC. Terraced houses are also relatively efficient, likely reflecting their usually modest size and the comparatively small exposed surface area of a typical mid-terrace home.

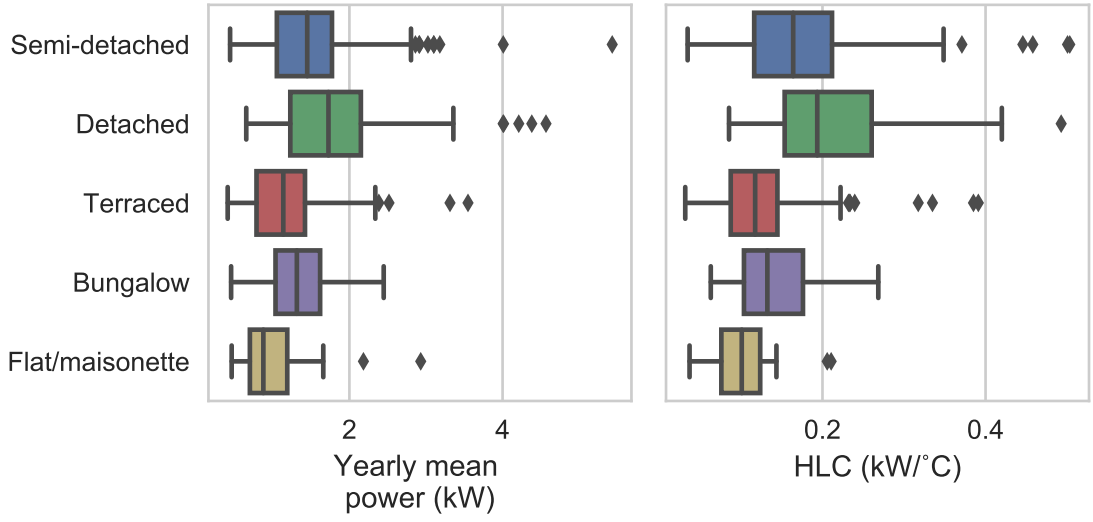

Figure 7.17: Comparison of mean yearly power demand and inferred HLC for EDRP-EDF sites as a function of dwelling type category

Although the distributions appear distinct, there is still considerable overlap between them. The dwellings can be further grouped according to dwelling size. Since surface area is not included in the metadata, the number of bedrooms is used as a proxy for size, following Summerfield, Oreszczyn, et al. (2015a). Bedrooms were chosen rather than total number of rooms because this latter value can be subject to some ambiguity in reporting, for example with regards to how partially separated dwelling areas should be counted (e.g. combined kitchen/dining rooms).

Figure 7.18 shows the HLC values grouped by dwelling type and number of bedrooms, demonstrating a clear effect related to dwelling size on top of dwelling type. Detached, semi-detached houses, and bungalows have the highest heat loss rates across the board, likely driven by their comparatively large exposed areas.

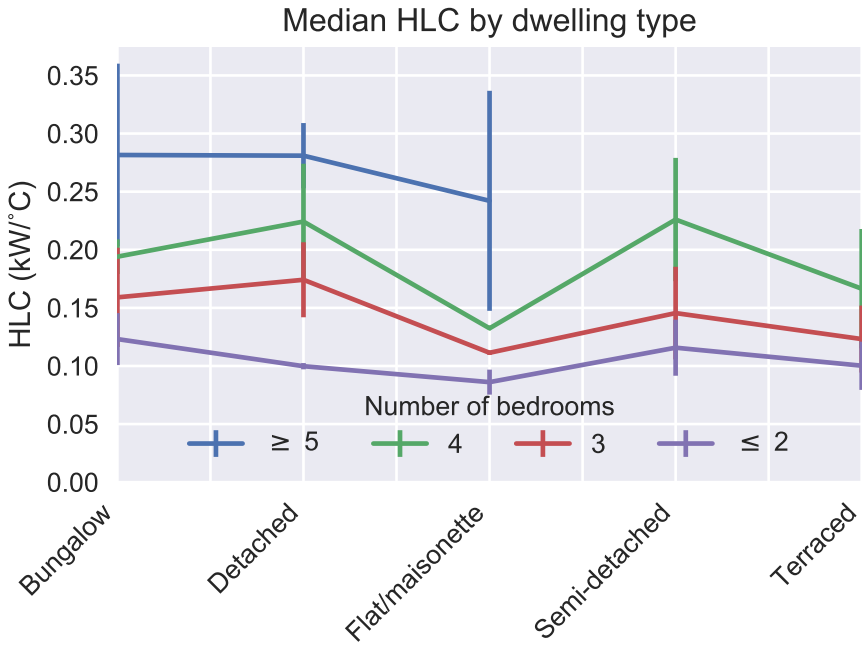

Figure 7.18: Median HLC for EDRP-EDF dwellings as a function of dwelling type and number of bedrooms, with error bars indicating the lower and upper quartiles

Figure 7.19 shows the power and HLC distributions across age bands, while Figure 7.20 shows the HLC values grouped by dwelling age and dwelling size. The trend by age is much smaller, notably for smaller dwellings of 2 or fewer bedrooms. Conversely, large dwellings of 5 or more bedrooms display the most marked improvement in performance for newer buildings, likely associated with improved construction standards.

7.3.2 Dwelling classification

Following the work performed on the Pennyland sites, these sites were classified using a GMM on the HLC values, with the aim of determining if the dwelling types could be distinguished by the classifier alone. Firstly, the significance of the effect of dwelling type on HLC was tested using

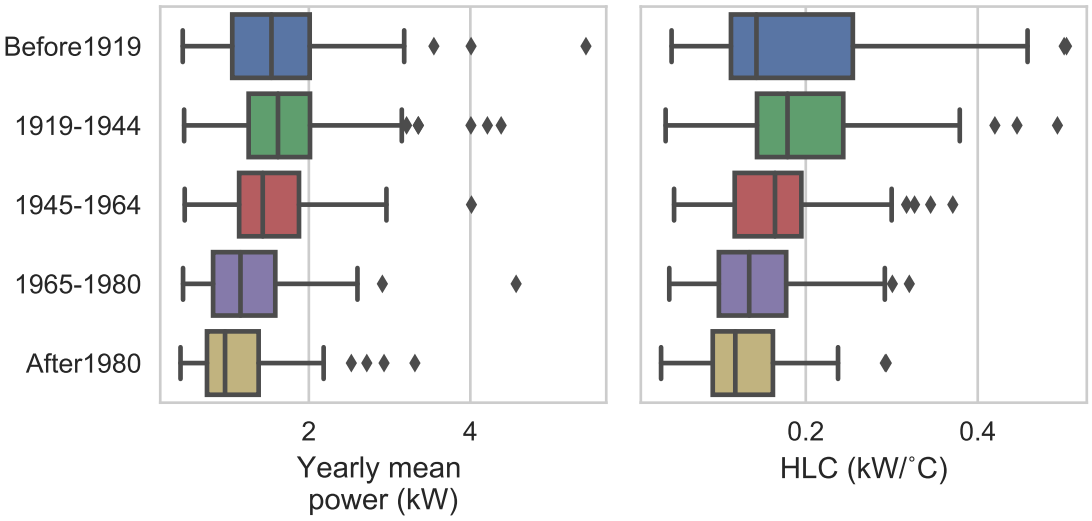

Figure 7.19: Comparison of mean yearly power demand and inferred HLC for EDRP-EDF sites as a function of dwelling age band

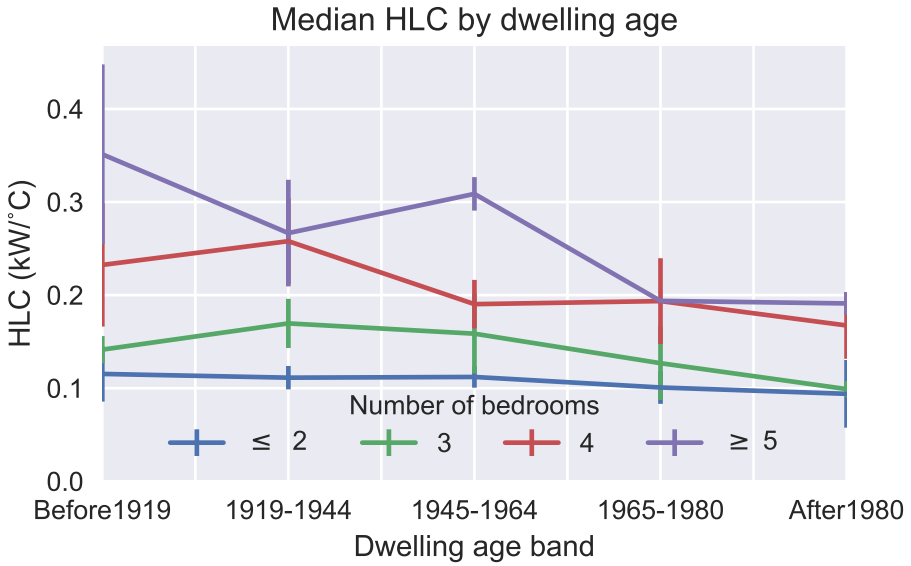

Figure 7.20: Median HLC for EDRP-EDF dwellings as a function of dwelling age and number of bedrooms, with error bars indicating the lower and upper quartiles

multi-linear Ordinary Least Squares (OLS) regression with replacement of dummy variables for each category. Coefficient p-values shown in Table 7.12 indicated that only the flat/maisonette and the terraced categories are significant, confirming expectations from Figure 7.17.

Table 7.12: Coefficient p-values for OLS regression of HLC against house type category

	Coef. p-value
Detached	0.624
Flat/maisonette	0.003
Semi-detached	0.231
Terraced	0.037

Three groups were defined following the OLS results: 'house', 'terrace', and 'flat'. The 'house' category combined 'Detached', 'Semi-detached', and 'Bungalow' dwellings. The GMM was allowed to classify dwellings into 3 unlabelled groups (group 0, 1, 2). The separation of sites into GMM groups is shown in Figure 7.21, which indicates that there is a poor correspondence between dwelling type and GMM group. The dwellings of a given type were not assigned to a given GMM group, instead they were spread across different GMM groups. There is extensive overlap between the HLC distributions, as well as a very small number of dwellings in the 'flat' category - this can be seen clearly in Figure 7.22.

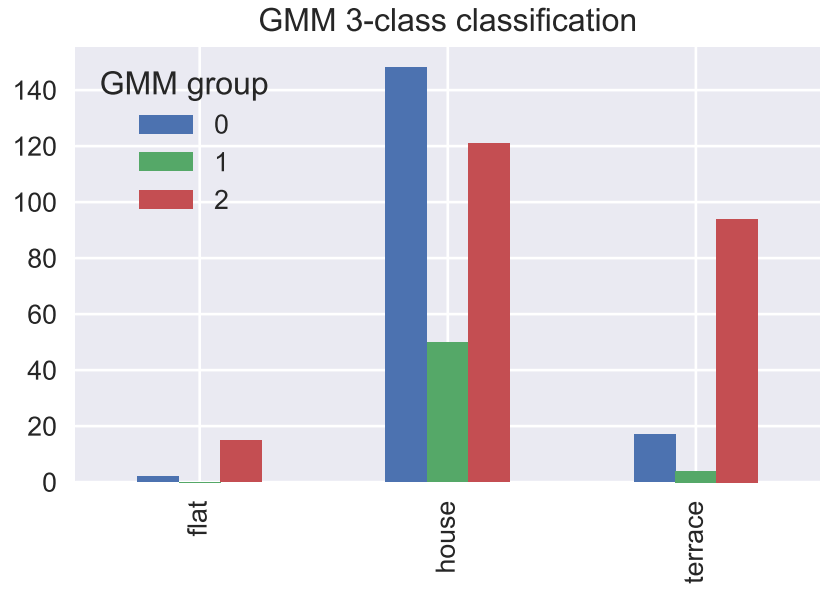

Figure 7.21: GMM classification on HLC estimates results for EDRP-EDF sites showing the number of sites of each known group assigned to each GMM group.

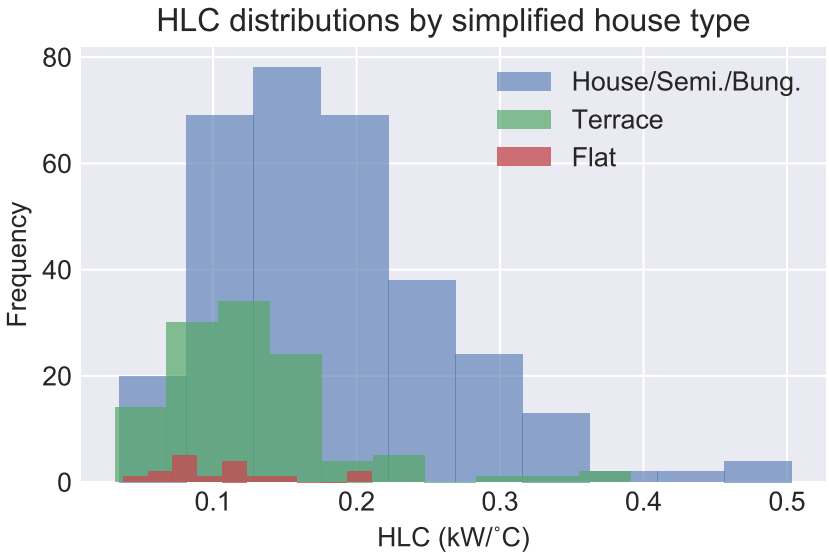

Figure 7.22: Overlaid distributions of EDF site HLC estimates grouped by simplified dwelling type (house, terrace, or flat)

In order to control for dwelling size, the 3-class GMM was applied to 2-bedroom properties only. This number of rooms was chosen for having a minimum of 5 dwellings in each category (Figure 7.23). However, the results (Figure 7.24) indicate that controlling for dwelling size does not separate the HLC distributions sufficiently to enable classes to be identified by a GMM classifier, as the dwellings are still assigned randomly to different GMM groups.

On the basis of the results from the Pennyland project, it was decided to simplify classes to 'flat' and 'house', with the intention of ensuring that the difference between the groups be large enough for the GMM to separate them out. The results in Table 7.13 show that flat classification is good,

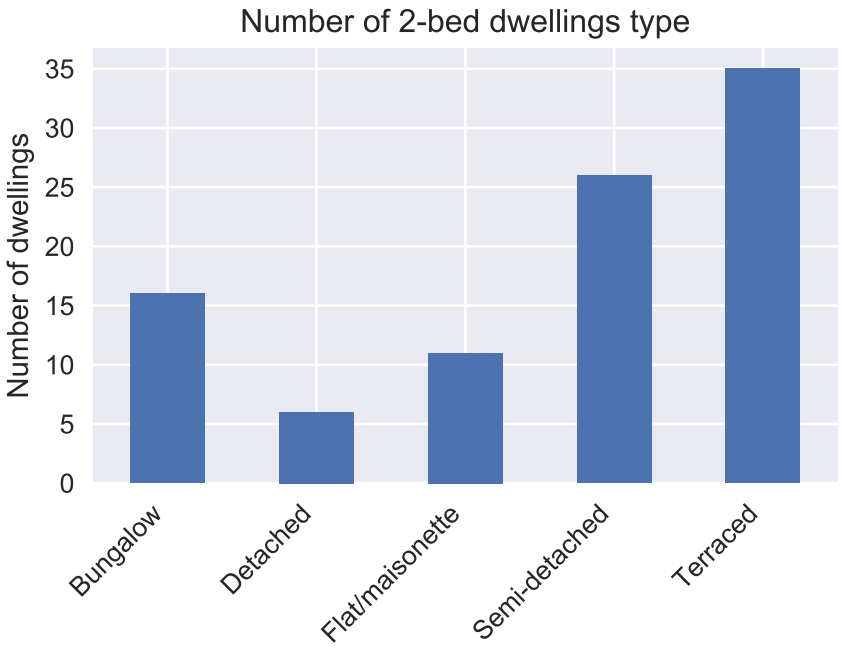

Figure 7.23: Number of 2-bed dwellings for each dwelling type

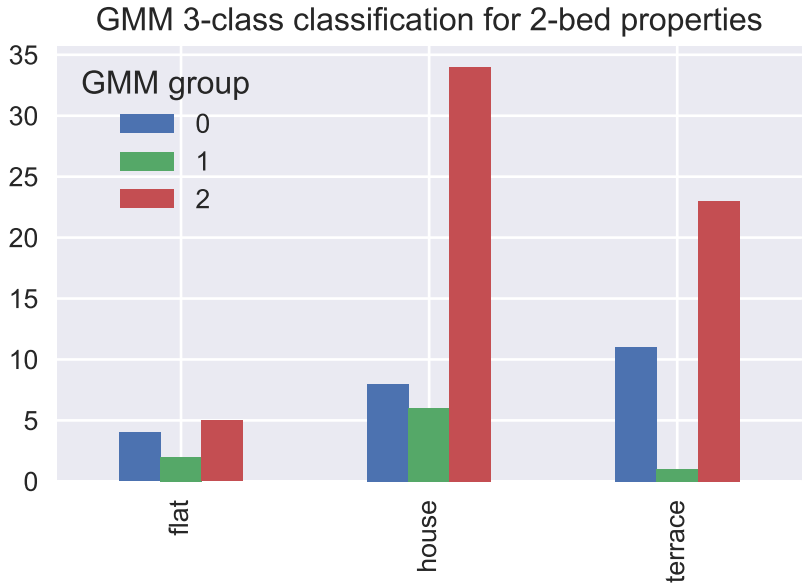

Figure 7.24: GMM classification on HLC estimates results for EDRP-EDF 2-bedroom properties only showing the number of sites of each known group assigned to each GMM group.

although there are very few sites, while the house classification is very poor. This underlines the finding that GMM cannot separate the HLC distributions if these are overlapping. We do not expect that other classification algorithms would fare any better - the distributions simply overlap too much to allow separation of dwelling types using HLC alone. Note that similar results were found when restricting the groups to only dwellings with 2 bedrooms.

Table 7.13: 2-class GMM classification results

	Flat	House
GMM group 0	17	249
GMM group 1	0	70
% Mis-Classified	0	78

7.4 Summary

This chapter presented three case studies applying the Deconstruct method.

The first demonstrated that Deconstruct can successfully be applied to a large dataset using state-of-the-art computing tools, and used this to compare HLC and NAC distributions for the EDRP dataset against EPCs and the CHM. This comparison indicated that EPC ratings tend to overestimate annual consumption, although the representativeness of the EDRP dataset remains questionable. The T_0 distribution for EDRP was also compared against the EFUS dataset nationally representative internal temperatures measurement. It was found that the distribution medians and shapes were not a good match, indicating that T_0 is not a reliable indicator of internal temperature. T_0 was particularly sensitive to the selection of the base-load gain parameter.

The second case study made use of a Gaussian Mixture Model (GMM) classifier to determine whether the HLC could be used to place dwellings into classes based on their thermal efficiency, placing them in low or high efficiency groups. This was performed using the Pennyland dataset. This was found to be successful provided that the dwelling types were highly distinct - when three groups were used to account for two variants of the high efficiency dwellings the results became unstable. Generated group means were found to be in good agreement with true group means. Performance of the classifier was found to be similar when using HLC and mean yearly power.

The third case study used the dwelling metadata of the EDRP-EDF dataset to explore the determinants of HLC values in terms of dwelling age, type, and size. The efficiency trends matched expectations, with efficiency improving from oldest to newest, from largest to smallest, and flats being the more efficient compared to detached dwellings. A similar approach to the Pennyland case study was used to apply a GMM to classify dwellings into groups, with the aim of identifying the dwelling class from its HLC value alone. However this was not found to be successful with this more heterogenous dataset. It was not possible to identify dwelling type from efficiency, even when dwelling types were reduced to only two categories and dwelling size was controlled for. This was due to the large range of HLC values in each category.

Chapter 8

Discussion

This thesis has presented the development and evaluation of the Deconstruct method, which aims to be a rapid, scalable method of thermal characterisation for UK dwellings using smart meter data. It has presented the model, the method of model parameter inference, the datasets used and the evaluation of the accuracy and relevance of the results achieved through applying the Deconstruct method to these datasets. This chapter is divided into sections following the aims and objectives outlined in Chapter 1 Section 1.4. Findings and limitations are summarised and discussed. The context and contribution of the research are also considered.

8.1 Dwelling Model

This research aimed to first determine the parameters needed for dwelling thermal characterisation. This section addresses that aim. The model theory and assumptions made for the thermal and internal temperature models are discussed, including the limitations of, and exclusions made from, the research scope.

The modelling philosophy adopted for the development of the Deconstruct method was to use physical theory and empirical simplifications to produce a model which could readily make use of additional data if provided. This was reflected in the development of flexible model code that allowed the use of assumed values or real data when available. It was developed in two stages - firstly a definition of the thermal balance of the dwelling based on results from literature was established, and secondly a system of linking functions to describe the relation between measured energy demand and thermal flows was developed. These functions provided an operationalisation of the thermal performance through the metered energy use.

8.1.1 Thermal model theory

The thermal model defined on the basis of literature described conductive and radiative heat exchange, thermal losses through ventilation, and gains from base-load power, as presented in Section 3.10 (eq. 3.25). The thermal model was further simplified in order to be applied to metered energy demand, and avoid the need for internal temperature monitoring. The Heat Loss Coefficient (HLC) was introduced to combine HTC and η_{HS}, as with metered consumption the heating system efficiency cannot be inferred independently. A simple linear internal temperature model was introduced (discussed in Section 8.1.3), parametrised by the internal/external temperature covariance C_T and the expected mean internal temperature at 5 °C T_0.

The total steady-state dwelling energy demand was given by eq. 8.1, the parameters of this model were summarised in Table 4.1:

$$P_{tot} = \begin{cases} \frac{1}{\eta_{HS}}|\Phi_{tot}| + P_B \text{ if } \Phi_{tot} < 0 \\ P_B \text{ otherwise} \end{cases} \tag{8.1}$$

Where

$$\Phi_{tot} = HLC((1 - C_T)T_{ex} - T_0 + C_T T_{fix}) + \frac{A_{sol}}{\eta_{HS}} I_{sol} + \frac{\eta_B}{\eta_{HS}} P_B \tag{8.2}$$

During the heating regime where $\Phi_{tot} < 0$, the total power was given by:

$$\begin{aligned} P_{tot} = &-HLC((1 - C_T)T_{ex} - T_0 + C_T T_{fix}) \\ &- \frac{A_{sol}}{\eta_{HS}} I_{sol} - \frac{\eta_B}{\eta_{HS}} P_B + P_B \end{aligned} \tag{8.3}$$

In order to apply the model, first simplifying assumptions for η_B and η_{HS} are made - although assumptions η_{HS} do not affect the HLC estimate. A value for C_T must be derived using internal temperature measurements. It was found that using nationally representative sample of dwelling monitoring produced an acceptable approximation.

For each site, one year of daily average total power demand P_{tot} was required, which may be calculated as the sum of P_{elec} and P_{gas}. It may be necessary to convert measured gas volume consumption to power using eq. 3.15. Corresponding daily average values for T_{ex} and I_{sol} were needed, these may be taken from gridded weather data using site location.

The base-load power P_B could be estimated as the average of sample of this data requiring high daily solar gains, while HLC and T_0 could be inferred from a low solar gain sample and A_sol from a low temperature sample. In all cases error, outlier, and where relevant absence filters must be applied to the input data. For each sub-sample a minimum of 5 daily average data points were

required after filtering, while overall sites were required to have no more than 50% of missing days of data. These requirements are no very stringent, making it possible to work with low quality input data.

This model is considerably more parsimonious than the SAP/BREDEM building model. An RdSAP assessment requires several pages of dwelling inputs to produce its estimates and generate an EPC. The principal tradeoff this model makes is requiring a full year's worth of daily total energy data. However, collecting such data is trivial with smart meters, which by default store 13 months of data in the UK. Barriers to obtaining this data are mainly related to data protection and data access issues rather than technical problems. Additionally, the subjective aspect of the RdSAP assessment by an inspector is avoided.

The HLC was found to depend on the internal temperature model approximation for C_T, but was robust to changes in η_B, P_B and T_0. η_B and η_{HS} could not be estimated directly from power-temperature data and had to be assumed or estimated through other means. T_0 and P_B were also estimated from data but are of secondary importance and are less robust.

The model was not found to be universal - for example on the EDF dataset with good input data, 18% of sites either had CVRMSE $> 30\%$ or HLC uncertainty greater than 30%, suggesting that in these cases the model did not suitably describe the data. It would be interesting for future research to further investigate what caused these sites to diverge from the thermal model. Initial manual investigation of some of these sites indicated that they simply displayed very high scatter resulting in high CVRMSE, while still broadly following the same linear demand trend. This may be an indication of large behavioural effects in the total power, for example if occupants manually control heating instead of using automated controls. It would be interesting to perform a more detailed study of a set of dwellings that display these properties, which could help determine to what extent occupant behaviour might be the cause.

It may be that for sites where the model fit is poor(CVRMSE is over 30%) an extension of the thermal model is needed. Although the fundamental physics which lead to the current model appear sound, there are additional refinements which could be considered. For example there is some evidence of 'droop' in demand at very low external temperatures, likely caused by undersized radiators relative to dwelling heat loss rate (note that it is unlikely that this is caused by undersized boilers, as the evidence from literature suggests they are usually oversized (HMGovernment 2010)). This 'droop' could be modelled with additional parameters such as a heating system maximum power. However, experience gained when analysing the data in this thesis suggests it will be difficult to establish with confidence using a steady-state model that an extended model with additional parameters is better than the simple one.

8.1.2 Heating system efficiency

The relation between HLC and HTC was explored in Section 6.4.2. This relation is important as the HTC is the value estimated by the co-heating test, which is probably the best established method for whole heating thermal performance measurement. The findings indicated that dwellings rank similarly whether their performance is indexed by HLC (which incorporates heating efficiency) or HTC (using reported site boiler efficiency). However, in some cases the boiler efficiency introduced a large discrepancy between these two values. Nonetheless, when thinking about overall dwelling performance, it is quite reasonable to consider the efficiency of the heating system to form an integral part of the assessment. In fact, EPC ratings do this implicitly, as they are based on yearly energy demand calculated using RdSAP, which takes heating system efficiency into account. This suggests that using the HLC would enable a reasonable comparison of performance between dwellings, without the need to determining the heating system efficiency.

Sources of data on boiler efficiency on a per-dwelling basis are limited. As a result, values from lab tests (such as SEDBUK ratings) or from various in-situ tests must be used to provide approximations, which introduces a new source of uncertainty. There is the possibility that smart boilers which record and report the readouts from a range of internal sensors become more commonplace. Originally these readouts were intended to facilitate troubleshooting but could equally provide boiler efficiency values. Heat pumps are one case where such data would be crucial. In traditional boilers the rated efficiency ranges from a minimum around 70% to a maximum in the mid 90% range - a 20% min-max shift and a much smaller inter-boiler range (assuming the distribution of boiler efficiencies is roughly Gaussian in this range). As a result the metered energy consumption is still relatively close to the supplied thermal energy, and the assumptions of the thermal model hold. Heat pumps however have a Coefficient of Performance (COP) of 2 or more, resulting in the delivered thermal energy being significantly different from the metered energy supply. Heat pump COPs are also external temperature dependent, further complicating matters. It would be expected to see a large difference between HLC and HTC for heat-pump equipped dwellings. Additionally, because of the temperature dependence of heat pump performance, the simple linear assumptions of the thermal model may break down.

An interesting case would be dwellings equipped with electric resistance heaters. Here, the heating system efficiency would be approximately 1 and the HLC and HTC would theoretically be equal. It was not possible to study such cases in this thesis, as the dwelling metadata did not make it possible to establish with confidence that a dwelling had an electricity supply but no gas - it was not possible to distinguish between dwellings which had no gas data because no gas supply was present, and those which had no gas data simply because it had not been recorded. The fact that British dwellings may purchase electricity and gas from different retailers makes such record keeping more complicated.

An important limitation of this method is its reliance on measured energy consumption. This

excludes all non-metered heat sources such as fireplaces and stoves, which remain common though are generally used infrequently. Solar thermal heating would also result in a more complex relation between metered demand and the thermal balance, particularly if solar hot water energy can be stored across multiple days. In this case the low solar irradiance days simplification would no longer be sufficient to factor out solar gains over time. The method furthermore requires that heat sources be metered at daily frequency or better. Oil burning boilers for example (still common in rural areas) can be metered but are not generally equipped for daily measurements. Similarly, if dwellings have on-site energy generation such as photovoltaic panels, and the power generated is not metered, this would also break the model assumptions.

8.1.3 Internal temperatures

In addition to the thermal model, it proved necessary to model internal temperatures. Initially, mean internal temperatures were modelled as being constant, consistent with implicit assumptions in the PRISM and PTG model approaches. However, theoretical treatment demonstrated that as a matter of basic physics, the mean daily internal temperature will depend on the external temperature (except in the case the heating system operating 100% of the time), independently of occupants choosing to heat to a different temperature - a fact already recognised in SAP (BRE 2014). This introduces a systematic bias if internal temperatures were assumed to be constant.

A simple linear model was proposed with two parameters - covariance of internal on external temperature C_T and reference internal temperature T_0, which is the expected mean internal temperature at 5 °C. This reference temperature was chosen to correspond with previous work in modelling internal temperatures - using internal expected temperature at 5 °C. Note that this linear model applies only to days which meet the low solar gain criteria used for estimating HLC.

When the internal temperature model was substituted into the thermal model, C_T was shown to be collinear with HLC. From a thermal physics perspective, the covariance C_T of internal with external temperature should itself be a function of the building heat loss rate, however the collinearity of the two factors made it difficult to solve this. Instead, the model was fitted to real temperature data and the suitability of using a simple UK mean which would address the systematic bias in the HLC for all sites was evaluated. The EFUS dataset provided a convenient nationally representative set of measured dwelling internal temperatures. Using EFUS, it was found that:

- The linear approximation was reasonable and that higher order models were not necessary.
- The C_T parameter could be estimated as well as its distribution using the low solar gains sampling method developed initially for estimating the HLC.
- C_T was negative for 6% of dwellings. These results were discarded as being non-physically consistent but, similarly to the 'failed' sites from the HLC estimate, it would be interesting

to further investigate what the cause of this relation is. For cases that were inspected manually, it was largely a question of large scatter in the temperature data resulting in greater uncertainty in the gradient. There may be however a deeper cause for some sites, such as systematic occupant responses to external temperature (increasing internal temperatures as external ones decreased).

By cross referencing EFUS with EHS it was furthermore possible to evaluate the dependence of C_T on other dwelling properties. It was decided based on existing research (Huebner et al. 2015; Stone et al. 2014) and exploratory analysis that only the physical properties of the dwelling need be considered - the dwelling age and dwelling type. No significant dependence was found on dwelling type, greatly simplifying further analysis as it avoided the need to consider dwelling type as an input to any further calculations. Dwelling age was found to have a small but statistically significant effect on C_T, with the effect being most notable for pre-1919 dwellings.

Some of the difference may be due to older dwellings not being heated in every room due to constraints imposed by the effectiveness of the heating systems, costs, or occupants' habits. To some extent, it is also related to the HLC of older dwellings being higher, resulting in a faster cool-down during non-heating periods when external temperatures are lower. As a result, it was decided not to adjust C_T for dwelling age when calculating HLC, since there would be a risk of a circular dependence - older dwellings would have a higher HLC, which in turn would result in a higher C_T. Using a dwelling-age adjusted value for C_T would be making an indirect assumption about the value of HLC. It was instead preferred to use the weighted average over all England sites of 0.17. The possible C_T distribution was instead incorporated into the uncertainty estimate (Section 8.3).

A possible avenue for future work would be to replace C_T with a model that calculated internal temperature taking HLC into account, for example by using parameters describing the heating pattern and dwelling cool-down curve. This would avoid any circular model parameter dependency issue. The Deconstruct method parameter inference method would need to be improved to allow the HLC and internal temperature model parameters to be estimated simultaneously.

An interesting possibility emerged from developing the internal temperature model which was to apply it even in the case that a dwelling had measured internal temperatures already. In previous research, heat loss rates have generally been calculated using the internal-external measured temperature difference. However, there are well documented difficulties in reliably measuring internal temperatures. Notably, measurements can only be taken in a few locations, usually in a subset of rooms. Additionally temperature measurements frequently have gaps and errors (this was observed with the SWI sites). Fitting the linear temperature model addresses these issues by effectively smoothing the input data and providing a model for filling gaps in the data. As long as enough temperature measurements are taken at a site for a linear regression to be performed, days without measurement data can be filled using the linear model. It also makes the detection of

failures and outliers simpler as these erroneous data points can be compared against the expected trend line. From the SWI results (Figure 6.25) this approach produces results as good as using measured temperatures directly.

The reference temperature T_0 was estimated per-site along with the HLC according to eq. 4.4. T_0 estimates were found to have comparatively large scatter compared to measured internal temperatures. At a population level, linear internal temperature models were fit to the EFUS dataset to generate a national distribution for T_0. This was compared to T_0 estimated from power-external temperature data in EDRP and found to be a poor match, with a particularly fat right-hand tail - i.e. a large number of dwellings with very high T_0 compared to EFUS. This was in part due to the dependence of T_0 on both the base-load and the base-load gains η_B. The original assumption that all energy use in the dwelling would eventually dissipate as heat does not appear to be valid. It may be possible to calculate a distribution for η_B by calibrating η_B and T_0 against measured internal temperatures for a nationally representative dataset which included both total power and internal temperatures.

Hot water usage is likely to carry significant amounts of energy directly out of the building through the drains. A rough estimate can be made of the hot water loss fraction using the median $P_B = 0.41$ estimated from EDRP and UK mean household hot water consumption of 16.8MJ/day = 0.19kW (Energy Saving Trust (2008)), if all of this energy was lost this would represent 54% of base-load energy loss, translating into $\eta_B = 0.46$. However the true value is likely to vary between dwellings, which could shift the T_0 estimate by 3 $^\circ$C on average and as much as 10 $^\circ$C (Figure 6.46).

Unfortunately, there are very few datasets with the data needed to answer some of the most interesting questions. There is in general a lack of data collection exercises that collect both total power and temperature data in every room.

EHS/EFUS cover a large part of the data collection needed and are particularly interesting for being nationally representative when paired with suitable weightings. However, they only include yearly electricity demand and therefore lack the necessary inputs for thermal modelling. An ideal development would be for future revisions of EFUS to include smart meter data and temperature in every room. With such a dataset it would be easier to determine the relationship between T_0 and true internal temperature, and additionally to better understand the base-load gains η_B.

While using measured internal temperatures clearly delivers more accurate HLC estimates, there is equally a clear benefit to using modelled internal temperatures as this enables the profiling of a much larger number of dwellings without additional sensors. The HLCs obtained presented good levels of agreement with measured internal temperature HLC despite the simplicity of the model. Furthermore, measuring internal temperatures presents a separate set of challenges, for example if a single temperature point is measured this may not represent the true mean dwelling temperature.

Smart thermostats are a promising development which could provide additional relevant data.

These would provide measured and set point temperatures as well as actual heating times. One advantage of thermostat measurements in this case is that they adjust the heating system according to their own measured temperature, so there must be a systematic relation between heat delivery and the thermostat registered temperature.

The simple temperature model makes a number of fairly strong assumptions. Notably, it assumes that the underlying daily temperature pattern does not change significantly over the days selected for the regression sample. Some causes of temperature change such as occupant absences during holidays should already be accounted for by the data sampling. However there could be significant changes to the heating schedule related to changes in occupancy, for example in family holidays, added occupants (guests/relatives) result in more rooms in use. Similarly there could be significant changes to zoned heating patterns resulting in large changes to the heat distribution and resulting heat transfer rates. This could also be related to changes in occupancy, for example heating up a spare bedroom. It is theoretically possible to detect occupancy changes from appliance energy use, although this is outside the scope of this thesis. Doing so would, however, require a significant labelled dataset to train any kind of algorithm. Ultimately, given the relatively large uncertainties on the relevant parameters, it is possible that attempts to improve model sophistication to take these effects into account would not lead to substantially reduced uncertainties in the results.

8.1.4 Limitations and exclusions

Firstly, as mentioned above, no attempt was made to construct an occupancy model, as this was judged to be outside of the scope of this thesis. In theory, it should be possible to develop a method of predicting occupancy level based on energy consumption data, for example by looking for traces of appliance usage in electricity data. However without a good labelled dataset this is difficult to do with any rigour. It was found that applying the simple lower bound cutoff heuristic described in the method worked well - although there is scope for future research.

One impact of not modelling occupancy is that metabolic gains are not accounted for. In general, these gains are small relative to the energy demand (in the range of 100W compared to power in the kW range), but nonetheless present an unknown source of error. Smart thermostats could again provide valuable additional information in this regard, as they are usually designed to adapt heating patterns to occupancy and do so by using a range of methods to predict occupancy, such as drawing on data from smartphone apps.

Secondly, although solar gains are described in the model theory, solar aperture parameters were not investigated. The use of a low solar gain sample alleviated the need to account for solar gains, which was identified as a significant effort in previous work (Stamp 2015; Jack 2015). The resulting HLC estimates were furthermore judged to be themselves among the best measures of thermal performance - solar gains were therefore set aside for future work.

This research has been resolutely on the physical side of the energy demand equation - as defined in the Method (Chapter 4). Thermal performance is operationalised through the physical relation between weather, heat transfer, and power demand rather than through the provision of energy service (thermal comfort) to occupants. Thermal comfort has intentionally not been considered as part of this analysis to avoid the increased complexity introduced by this higher level concept - there are enough challenges in modelling the purely physical side of energy demand.

8.2 Data sources and parameter inference

8.2.1 Deconstruct data requirements

The final model requires total power, external air temperature, and solar irradiance. Although wind speed and precipitation were used during the method development, they were not found to to be significant for the final parameter estimates. Daily averages of power and weather were found to be suitable, with no advantage to averaging over longer periods. Rolling/window averaging was considered to reduce scatter, but this would undermine the independence of days which enables arbitrary subsets to be selected. Total power was defined as electricity plus gas, although this could just be electricity in the case of electric heating. The datasets provided did not reliably identify dwellings with electric heating (e.g. EDF has a single or dual tariff identifier field but does not specify a whether single tariff electricity customer might be getting gas from a different supplier). Analysis could be expanded to more dwellings if this information was known.

One of the key aspects of the Deconstruct method developed in this research was a 'natural experiment' approach of discovering ideal samples in historical energy data suitable for estimating given parameters. This was found to be successful compared to attempting to fit the complete thermal model to data - the latter approach required excessive manual tweaking of solver parameters, which is unacceptable for a scalable approach to thermal assessment. The chosen sampling strategy delivered a method robust to poor quality and noisy input data. The improvement of the model fit is significant as shown through the reduction in CVRMSE. The cut-off values for the various parameters were chosen in line with recommendations from literature where available, however a familiarity with distributions of weather and energy parameters would be needed to develop an intuition of the expected range and suitability of cut-offs. More extensive work could be done to deepen the understanding of the effect of these parameters.

At least one year's worth of data was required for each site in this analysis, which is more than strictly necessary, given that for the HLC estimate only 5 or more daily data points are required. Requiring a year's worth of data does limit how quickly the method can be applied - it certainly takes much longer than the 2-3 weeks generally used for a co-heating test. However the tradeoff is in robustness (notably to solar gain induced errors), in simplicity of application, and in not

requiring internal temperature measurements. Requiring a full year was found to be the best method to offset the problem of bad inputs, since sites with less than a year of data in experiments which range over at least 2 years were generally those with severe data collection problems.

8.2.2 Dataset outputs

In addition to the Deconstruct method, the datasets described are themselves important outputs of this research. Cleaned and weather-referenced versions of the raw datasets represent a considerable investment of time, effort, and technical know-how. The output datasets include:

- 'Re-mastered' Solid Wall Insulation (SWI) trial data. The SWI dataset required extensive work solving data import issues and running data quality assurance. A heavily cleaned version was finally produced which included site monitoring data and site metadata painstakingly extracted from the site report PDF files.
- Cleaned versions of the EDRP-EDF dataset. The raw input files had been transmitted and format shifted a number of times resulting in odd glitches. A clean and error-free set of tables with a more practical column layout was produced.
- The full EDRP was already cleaned, but was delivered in a relatively inconvenient format. Importing it into the SQL database used for most of the data work enabled convenient re-exports for a variety of purposes.
- MetOffice weather data on a regular longitude-latitude grid.

These datasets were shared where permissible with other researchers, as well as with collaborators at EDF R&D UK.

8.2.3 Energy data access challenges

One of the most significant challenges with regards to data in this research was not processing large datasets (which was covered in the Datasets chapter) but obtaining them and dealing with anonymised data. UK law requires privacy to be protected under the Data Protection Act, with the broad directive that personally identifiable information should not be released publicly. However there is a lack of systematic guidelines for the handling of energy data, and procedures that would enable access to original data where necessary. Anecdotally, it is easier to obtain highly sensitive medical record data for research purposes than it is to obtain home energy data. Energy data anonymisation approaches suffer from something of a 'cargo cult' approach - if enough 'anonymisation-like' transformations are applied, the resulting data will meet the data protection requirements. In some cases this may in fact be the case, but only inasmuch as all useful data has been removed - in which case a randomly generated dataset would provide as much relevant information.

The publicly released EDRP dataset was a particularly egregious example, where for many sites any location data more specific than the large NUTS1 administrative region was removed. The rational for this was to ensure the data was not 'personally identifiable' by avoiding associating energy data with an exact address, on the basis that the address is a personal identifier. However, addresses themselves are public domain. This might be justified if other dwelling information had been provided which could breach privacy if associated with an exact address (e.g. number of children in the household), however this too was removed and replaced with 'ACORN' codes which, as a commercial classification for the marketing industry with a non-public methodology and which in any case are drawn from postcodes rather than being dwelling specific. As a result these codes do not provide relevant information for this research.

Ironically, domestic EPC data was later released which includes energy consumption information and exact addresses. This would have been extremely interesting to cross-reference with EDRP smart meter data - however the raw EDRP dataset and its included identifiers have been deleted and only the anonymised version retained in the UK data service archives - a sad loss for domestic energy research.

It would be highly desirable for energy data to be made available for research in a way that does not remove valuable information. A reasonable access approval process for researchers (for example subject to ethics committee approval) would be preferable to full open access to heavily censored datasets. A more systematic understanding of the real privacy risks would help allay concerns regarding the sharing of such data. There is in the UK something of a mismatch between the privacy expectations regarding many aspects of daily life (such as the prevalence of CCTV) and the limitations imposed on energy data in the name of privacy. Public release of energy information may not be desirable or necessary, whereas release for research purposes would be beneficial on the understanding that a reasonable ethical approval process is put in place.

8.3 Uncertainty

This section addresses research question 2.1 "How precise, accurate and reliable is the method?".

The lack of alternative measurements of HLC made the estimation of the true error in inferred HLC difficult. The ideal solution would be to perform co-heating tests to obtain HTC for a representative set of dwellings and also obtain smart meter readings for those same dwellings when occupied, as well as monitoring of boiler performance so that the HLC and HTC values can be compared. While no such dataset currently exists, the possibility of evaluating in-situ performance assessment methods should encourage the aggregation of co-heating test results and the cross-referencing with dwelling smart meter data when the tested dwellings are occupied. Additionally, as smart meter data becomes more accessible via energy suppliers and energy data brokers it could be that smart meter data may be found corresponding to previously tested dwellings.

The Monte Carlo (MC) approach to parameter estimate uncertainty recommended by the JCGM produced very promising results, estimating the HLC uncertainty at 15% which is in-line with the work of Jack (2015) who performed co-heating tests on test dwellings with simulated occupant interactions. It also compares favourably with co-heating test uncertainties of 10% in unoccupied dwellings - this representing the best-case scenario. It should nonetheless be noted that with a 15% uncertainty, changes in a given dwelling's performance from retrofits need to be relatively large to be detected with confidence, which may be a problem when expected retrofit savings tend to be in the 20% range. This is still an improvement over comparing measured consumption against SAP predictions however, as with this method the same assessment can be performed before and after retrofit - an important improvement in test consistency.

The MC approach was a good fit with the modelling philosophy of building a model which can readily make use of additional data. Input distributions could be developed for every model parameter to refine the uncertainty estimate in each output variable. For example, the previously discussed assumptions on boiler efficiency could be substituted by an empirical efficiency distribution and the resulting effect on HLC discovered.

The MC approach could also be expanded upon through a full Bayesian analysis. This approach could make better output predictions by using known distributions for inputs and assumptions. Bayesian statistics were consciously avoided for this initial work for simplicity and to provide a clean baseline approach. Additionally, it was preferred that a method intended to act as an empirical alternative to the EPC should be relatively simple, in order to promote communication and dissemination of the method.

When including other approaches to error estimation, including year-on-year power prediction and comparison of modelled with measured internal temperatures, errors and uncertainties were consistently found to be in the 15-20% range, which while not small do compare favourably with the 10-15% uncertainties in HTC estimates.

8.4 Model evaluation and application

This section discusses the comparisons between the Deconstruct method outputs and other methods of calculating dwelling thermal losses such as the use of intrusive monitoring and EPC assessments.

8.4.1 HLC using monitored internal temperature

The main form of internal monitoring data available was measured internal temperatures which were used to contrast HLC estimates made using the internal temperature model against those made using measured ΔT. Only the SWI dataset provided both the daily temperature monitoring

and energy data needed to perform the comparison - and this data was of generally low quality. The positive side is that this represents something of a worst-case scenario - poor quality data from old dwellings. In this context, the mean bias-adjusted error of 18.9% (Table 6.3) compares favourably with the 15% predicted uncertainty from the MC approach. The choice of C_T has a direct scaling effect and therefore is key in adjusting the bias. Although C_T was shown to depend on dwelling age, the results from the EDF dataset indicated that this has no practical impact and therefore the dwelling age adjustment can be ignored.

In general, the mean C_T value used acts as an empirical correction to address a source of systematic bias in the HLC estimate. If the same value is used for all sites, when the HLC is used as an index of performance this scaling would not have any impact on the relative performance ranking of dwellings.

8.4.2 Predicting power demand with Deconstruct

Comparing the HLC results for the same sites for subsequent years was used to establish the internal consistency of the model - a simple reliability metric being that the dwelling HLC should be constant over time. On average the HLC prediction error was 0±23%. This indicates a good performance in aggregate while per-site errors were of a similar order to those expected from the MC uncertainty approach. Using the HLC from one year to predict the heating power demand in the following year also yielded good results, with a 6% median bias and 15% inter-quartile range. These results support the basic theoretical basis of the Deconstruct method in demonstrating consistency in results.

A small number of sites in the power prediction experiment displayed very large changes in the HLC across years. These sites appeared to display significant shifts in energy demand patterns that could have been related to retrofits performed in between heating seasons. An interesting avenue for research would be to cross-reference dwelling Deconstruct analysis results over years with records of dwelling retrofits to investigate the impact of known works on the HLC values.

Calculating a Normalised Annual Consumption (NAC) enabled a comparison against similar existing estimates provided by the Cambridge Housing Model (CHM) and EPC database. These results demonstrated how in principle the Deconstruct method can be compared against existing methods, and also suggested that RdSAP based methods systematically over-estimate normalised consumption. However, in order to obtain the NAC an estimate of the effective solar aperture was needed - this parameter was not subject to in-depth study due to limits on the time and scope of this research. As a result there remains some uncertainty about the representativeness of the results.

Solar gains calculations are perhaps less important for dwelling performance evaluation but are nonetheless necessary to generate NACs, which in turn are necessary for comparisons with RdSAP

outputs (whether through EPC or CHM) because frequently only the yearly calculations are available for these. There are still important issues to address, such as the correct solar projection to use (e.g. horizontal or vertical plane) and the physical interpretation of the resulting parameter.

8.4.3 Scaling Deconstruct to large datasets

The case study performed in collaboration with the UKDS demonstrated the suitability of Deconstruct to be applied to large smart meter datasets (such as EDRP) using an industry-standard big data software platform (PySpark). The 63% success rate was somewhat lower than the EDRP-EDF subset, but should be contextualised by the paucity of location data in EDRP. Since this has a direct impact on the accuracy of the weather lookup, if exact address data had been available it is expected that the success rate would be closer to the 82% observed for the EDRP-EDF dataset. This project indicates that a majority of dwellings could be assessed with the Deconstruct method if suitable data is available. Furthermore, so long as smart meter data remains available for a dwelling the HLC could be updated every year and a time series of thermal performance generated for each dwelling. The fact that the application has already been demonstrated using state of the art large data processing tools suggests that this would be a feasible project. Furthermore, it is expected that when scaling to a larger number of dwellings, the manual aspects of the process would become a much smaller proportion of the total analysis load as all the main steps may be automated.

8.4.4 HLC by dwelling type groups

The estimated HLC values were also able to successfully classify dwellings from the Pennyland project into their correct high and low efficiency construction types, with a similar success rate to using yearly mean power demand. Broadening this analysis to the population of dwellings in the EDRP-EDF dataset it was found that HLC distributions overlap too much to allow correct automatic labelling of dwelling types. Further restricting the inputs to dwellings of similar size (2 bedrooms) did not improve the classification.

Despite this, when dwelling properties such as age and type are know, plots of HLC clearly reflected the difference in performance as a function of age and size that one would expect to see, with dwellings becoming more efficient over time and large dwellings having a larger heat loss that smaller ones. It would be interesting to compare dwellings HLC as a function of total floor area directly instead of using bedrooms as a proxy for dwelling size.

8.4.5 Evaluation limitations

It was not possible to assess the performance of the method to identify savings from retrofits. The original intention had been to use the SWI dataset for this purpose, however the poor data quality and lack of information about the retrofit procedures carried out made this impossible. Studying the change in HLC as a function of retrofits would be an extremely interesting project, as it would enable much better quality control, performance assessment, and cost effectiveness calculations for home improvements.

It was also not possible to establish beyond doubt that the HLC inferred using Deconstruct is independent of dwelling occupants. The theoretical underpinnings of the model are purely physical. From this it may be inferred that if the purely physical model is a good fit for observed data, then the known physical processes explained the observations and occupant behaviours are not needed to explain these observations. Furthermore, by choosing suitable data slices - notably the low solar gain sample - it is aimed to apply the model under conditions which minimise expected occupant perturbations of the thermal system (such as window opening). However, it was not possible to empirically test this hypothesis. To do so would require a more controlled trial. For example, HLC for dwellings monitored over a long period during which different occupants were resident there could be compared between different years to determine whether the occupant affects the HLC value. Achieving this would require close collaboration with landlords (this could be done for example through a housing association).

8.5 Alternative approaches

During the progress of this research, a number of approaches to dwelling thermal modelling were experimented with but were ultimately discarded.

8.5.1 SAP based models

It has been relatively common in UK dwelling energy research to build improvements to domestic energy models on the basis of the SAP/BRE model. This has shown some success and had led to revisions of the standard to be published every few years to take the model refinements into account. This route was not taken for this research, preference begin given to a data-driven empirical model. Such a model has the potential to be a better fit for a ‘data rich’ future energy system. Furthermore there are fundamental limitations to the approach of refining the SAP model with regards to the assessment of existing buildings through RdSAP. The need for a manual inspection introduces the problem of human error and implies a certain cost and inconvenience for inspections to be performed, at the end of which the model accuracy is unlikely to increase significantly.

8.5.2 Black box models

A range of models could be considered to fall into the 'black-box' category, in that they make use of more advanced statistical methods at the cost of physical interpretability. A conscious decision was made early on not to pursue black-box modelling methods, despite existing successes in predicting electricity demand with black-box machine learning algorithms such as neural networks. The fundamental problem with black-box models is their lack of interpretability. In some of their popular applications, such as image recognition, there is no interest or need to understand model coefficients, which in any case have no physical meaning - success is measured purely on prediction accuracy. However, dwelling thermal properties have a sound basis in physical theory. On the one hand, this enables the introduction of a number of simplifying assumptions about the model form. On the other hand, the physical properties of dwellings are of considerable interest when considering the performance of different types of buildings, building materials, or the impacts of changes and retrofits.

One empirical approach considered was using Multivariate Adaptive Regression Splines (MARS) models, a form of non-parametric regression that enables the fitting of discontinuous linear models (models formed of several linear parts with different parameters) (Friedman 1991). The interest of such an approach was to automatically find inflexion points, such as the heating on/off boundary which is a function of temperature and solar irradiance. However, this method generates model segments independently of any physical theory. As a result, interpretation of model results becomes difficult. The MARS algorithm produces a piecewise linear function, but with no guarantee that a model fit will produce a given number of model segments or that inflection points will occur in physically meaningful dimensions or locations. It was also found to generate spurious model structure at data boundaries. For example there are a relatively small number of days with extreme conditions (such as extreme cold), and these tend to have high scatter. The MARS algorithm sometimes, but not always, introduced a new model region to describe these. The relative obscurity of the MARS algorithm was also a concern due to the comparative shortage of literature on the subject. More popular polynomial regression models (commonly known as LOESS models) apart from being better documented, incurred similar drawbacks to MARS, with the further complication that physical interpretation of polynomial coefficients was exponentially more complicated.

Overall the fundamental issue was the lack of reliability in the algorithm outputs - given noisy energy data inputs there was no guarantee of being able to generate consistent models across many sites. While there is value in generating a model with a good fit for a single site using this algorithm, it makes it difficult to perform comparisons between dwellings.

8.5.3 Co-heating test based models

Regular co-heating tests are considered the state of the art for determining existing dwelling thermal performance. However, there did not appear to be any means to scale this approach to large numbers of dwellings, making it a non-starter from the scalability perspective.

Other research has taken a similar aim of adapting the principles of the co-heating test to occupied dwellings, notably the work of Jack (2015) - who performed co-heating analyses on four test dwellings equipped with an automation system which simulated occupant behaviour. This work focused on short (2-3 week) monitoring periods but required internal temperature monitoring (implying additional manual labour and equipment) as well as manual data processing.

8.5.4 Dynamic models

Dynamic thermal models have been identified as having a number of advantages over steady-state models for thermal modelling. During the course of this research, such models were duly considered to the extent of following an intensive course on the use of the dynamic models developed as part of IEA Annex 58. These models were ultimately not pursued for a number of reasons. The data quality issues previously mentioned made dynamic modelling difficult. From literature and experience it emerged that dynamic models require continuous data series free from significant outliers. The input datasets included large numbers of missing data points, therefore applying dynamic models would require extensive gap-filling and/or the creation of algorithms to find continuous data segments within the larger datasets - itself an error-prone process. These limitations would severely undermine the scalability goals of this research. Furthermore, dynamic models are an extension of steady-state ones. It was therefore logical to establish a working steady-state model before moving to dynamic ones.

Perhaps the most interesting aspect of dwelling demand that may be addressed by dynamic models is the dwelling thermal mass. As a steady-state method, Deconstruct does not address this. Thermal mass may be important in overall dwelling performance by acting as a buffer to changing conditions and storing energy across days, reducing energy demand. Dynamic analysis would be needed to better understand the impact on HLC estimates. Passive house designs would potentially be most significantly affected, as they are generally designed with very high insulation and thermal mass. As a result, power demand is unlikely to vary significantly with external temperature making HLC estimation difficult, while the thermal mass of the dwelling may further undermine the model assumptions.

8.6 Summary

This chapter discussed the key findings of the thesis. The research questions were addressed with reference to results presented previously. An appraisal was made of the developed models and methods. The data requirements for the Deconstruct method and challenges surrounding acquisition and processing of this data were discussed, as well as ideas for future data collection. The uncertainties in Deconstruct results were also summarised. The evaluation of the model and its application to relevant datasets was discussed, with the findings that HLC estimates appeared to be consistent and should be reliable within the expected uncertainty range, while other parameters, notably internal temperatures, could not be considered reliable. A number of alternative approaches that were considered but not used for this research were also discussed. Further validation work is recommended, notably in terms of comparing heat loss coefficient estimates from occupied dwellings against co-heating tests and further establishing that the HLC estimates are independent of occupants.

Chapter 9

Conclusion

This thesis has aimed to develop a method of characterising UK dwelling thermal performance which could readily be applied across millions of dwellings, taking advantage of the expanding availability of smart meter data. It established the context which motivates the development of new methods to infer dwelling performance from smart meter data in Chapter 2. A model of energy demand based on dwelling thermal physics was presented in Chapter 3. Chapter 4 discussed a novel 'natural experiment' approach and method for inferring model parameters from existing dwelling data. This method was given the name 'Deconstruct'. Chapter 5 presented the datasets used to develop and evaluate Deconstruct. Chapters 6 and 7 presented an evaluation of different model aspects, application case studies, and the results achieved. The findings were discussed in Chapter 8. This chapter summarises these findings and presents the broader impact and outlook of the work.

9.1 Findings

The key findings of this research are as follows.

- Using the Deconstruct method it was possible to infer a dwelling's Heat Loss Coefficient (HLC) with a mean uncertainty of 15% for 63% of the dwellings studied, using daily average smart meter data for one year and dwelling location, from which local temperature and solar irradiance was derived. This demonstrates the ability to non-intrusively estimate an indicator of thermal performance of a dwelling.

- The method was demonstrated to work on large datasets (over 7000 dwellings) using state-of-the art Big-Data processing infrastructure, supporting its potential as a scalable approach to estimating the thermal performance of smart meter-equipped UK dwellings.

- The method was to be able to separate dwellings into efficiency groups matching their known construction type, demonstrating the method's suitability for identifying dwelling physical characteristics.

- A nationally representative mean value for the internal-external dwelling temperature covariance parameter was found for the UK, $C_T = 0.17$. This value contributes to the calculation of the HLC, and was found to be a reasonable approximation for most dwellings. This was important as it enables HLC estimates without the need for internal temperature monitoring, allowing entirely non-intrusive HLC estimates.

The application of the method was found to be subject to the following constraints.

- One year of daily average total energy consumption and data is required, by which is implied electricity and gas readings with no more than 50% of missing data and no less than 5 data points in each sub-sample described in the Method Chapter 4 (high solar gain, low solar gain, low temperature). Daily average values for weather variables including air temperature and solar irradiance are required, these may be derived from building location.

- Location data is required with good precision to reduce HLC uncertainty - ideally full addresses which may present data protection challenges. Increasing the location specificity from NUTS regions in the EDRP dataset to postcode areas in the EDRP-EDF subset resulted in a corresponding increase in the percentage of sites for which HLC could be estimated from 63% to 82%.

- It was possible to estimate a value for the dwelling reference internal temperature T_0, but the sensitivity of this value to input assumptions, notably the base-load gain factor η_B, implies that this parameter is not a reliable indicator of dwelling internal temperature.

- It was not possible to infer dwelling heating efficiency separately from fabric heat loss. Therefore the Heat Loss Coefficient (HLC) thermal performance characteristic was defined, which combined Heat Transfer Coefficient (HTC) and heating system efficiency η_{HS}. It was demonstrated that dwelling performance ranked similarly whether HLC or HTC was used, indicating that HLC would work as well as HTC as an index of thermal performance, while only requiring smart meter data as an input.

9.2 Contribution

The key contribution of this research is the Deconstruct method, which enables non-intrusive estimates of dwelling energy demand characteristics, most notably the Heat Loss Coefficient, for dwellings equipped with smart meters.

Efforts to develop thermal performance assessments of occupied dwellings was highlighted both by

the interest of the UK government through BEIS (BEIS 2016c) and the launch of IEA Annex 71 (Roels 2017). Both highlighted that there are currently no established methods - the Deconstruct method proposed in this research contributes to this emerging field. There is, furthermore, a desire to develop an 'empirical EPC'. The HLC estimated through Deconstruct could contribute to such an empirical dwelling performance assessment.

In general, the Deconstruct method is robust and parsimonious in terms of the quality and number of data inputs, facilitating its application to dwelling energy monitoring data. The UK government plans to enable access to dwelling smart meter data through the Data Communications Company (DCC). This should enable entities such as utilities and research agencies to access dwelling energy data for large numbers of dwellings and to retrieve data for long periods of time for each dwelling. The DCC greatly facilitates collecting a full year of data for dwellings. The Deconstruct method is ideally placed to take advantage of this data source to provide thermal performance estimates for connected dwellings. This also enables additional value to be derived from meter data for utilities and occupants.

Furthermore, using such data sources would allow HLC estimates for connected dwellings to be updated each year. This would be relevant for supporting energy policy making at the national or regional level and tracking the effect of such policies over subsequent years. In this respect the relative simplicity of the Deconstruct method is an advantage as it can facilitate implementation and communication of results across domains of expertise.

For the purposes of dwelling retrofit assessment, Deconstruct offers a tool which could enable before and after heat loss assessments of dwellings undergoing retrofits. The HLC uncertainty of 15% implies that only relatively large changes in fabric and heating system efficiency will be detected, but this could nonetheless be used as a first order verification and diagnostic tool.

Other contributions of this research include:

- Data Extraction, Transform, and Load (ETL) platform code for the processing of the datasets described. This code facilitates reproducibility of the research.
- Cleaned and standardised dataset tables in the open Tabular Datapackage format (FrictionlessData (2017)), facilitating dissemination to other researchers in the field. This enables researchers to save considerable time, effort, and knowledge acquisition to produce cleaned datasets.
- Conference papers Chambers, Oreszczyn, and Shipworth (2015), Chambers et al. (2015), and 'EPC toolkit' software (Chambers and Stone 2016).

9.3 Future Work

This research indicated a number of directions for future work.

9.3.1 Improvements to the Deconstruct model

Some aspects of the Deconstruct model were not studied in depth, given limitations of time and scope imposed by the desire to focus on the HLC as the primary indicator of dwelling thermal performance. These aspects notably include solar gains through the effective solar aperture and solar radiative losses. Both of these introduce more complex geometric considerations to the model. It would be useful to determine to what extent they may improve the performance of the Deconstruct method in terms of parameter inference and power prediction accuracy. This could include for example a comparison of geometric projections of solar irradiance (e.g. on vertical planes with various orientations) as an input to the thermal model.

The base-load gain parameter η_B and the internal temperature T_0 were found to be closely linked by the thermal model. Obtaining a better understanding of base-load gains would help address a significant source of uncertainty in T_0. The results so far indicated that T_0 could not be considered a reliable index of internal temperature because of the potentially large inter-dwelling variation of η_B. If a method of estimating η_B could be developed, this could at least place clearer bounds on the range of η_B between dwellings and therefore on the uncertainty in T_0.

The heating system efficiency parameter η_{HS} is an important component of the thermal model, but one for which there is currently limited data on a per-dwelling basis. Internet-enabled heating systems (smart thermostats and smart boilers) are a current trend that may provide such relevant data. Deconstruct is structured in such a way that integration of such data should be straightforward. These systems could provide valuable temperature and set point time series, and possibly even report real time boiler efficiencies and occupant presence. These would demonstrate the significant added value of combining data from smart meter and smart device sources, which in turn could help spur further adoption of these technologies.

Dwelling occupancy is another possible area for further development. Occupancy was modelled in a simple way, using daily consumption cutoffs to identify unoccupied days. It could be interesting to develop more sophisticated models to determine whether a dwelling is occupied, for example using sub-daily energy demand patterns. This could be compared with the simple approach taken in this research to determine if there is any improvement in the HLC estimate.

Experience gained developing Deconstruct indicated that a Bayesian approach could be well suited to this problem. Bayesian statistics make it possible to take advantage of the existing research in various aspects of dwelling energy use by using the knowledge gained to inform the prior distributions for model fitting. This would enable a grey-box model approach that would take better advantage of the existing body of research. Bayesian methods were not used for this research as it was preferred to establish a basis for the Deconstruct method using widely accepted and understood statistical methods.

9.3.2 Representative study including internal temperatures

None of the smart energy meter datasets obtained for this thesis were designed to be representative of UK dwellings. This makes it harder to extrapolate the findings to other dwellings. Collection of such a dataset would therefore be very useful. It would additionally be important for the dwelling data to include internal temperature monitoring, which would enable a cross comparison between modelled and measured temperatures and their impact on inferred thermal parameters for a representative sample of dwellings. Performing such data collection would likely require a substantial research budget and was outside of the scope of this thesis.

One possible source for this data would be for a future version of EFUS to include total metered consumption using smart meters in addition to the internal temperature measurements. EFUS and EHS already collect all other data relevant to thermal modelling, and are nationally representative. Since smart meters are to be installed in most dwellings, it does not seem unreasonable that future iterations of EFUS would be able to collect their data in addition to the variables monitored previously.

Having access to measured internal temperature would present the considerable advantage over EDRP of enabling comparison between modelled and measured internal temperatures, and potentially enabling refinements to the model to enable accurate T_0 estimates yielding the practical benefits outlined previously. This would also address a general shortage of studies with both temperature and total energy (electricity and gas).

9.3.3 Comparing Deconstruct against other dwelling assessments

In this research it was possible to compare the distribution of Normalised Annual Consumption (NAC) (calculated using Deconstruct parameter estimates) against the national EPC distribution, but not to compare the NAC and EPC on a per-dwelling basis. The EPC database contains full addresses which would make it possible to match dwelling smart meter data against it, if the address of the meter is known. Alternatively, it may be possible to make use of the smart meter data access capabilities proposed by the DCC to obtain smart meter data that can be linked to EPCs. This could enable a systematic study of the relation between bottom-up and data-driven dwelling thermal performance estimates.

Co-heating is generally recognised as the state of the art in terms of whole-dwelling as-built thermal performance measurement. Comparing HLC with the co-heating HTC estimate would be a valuable exercise, but would require a co-heating test for dwellings which were subsequently occupied and for which smart meter data was collected. Coordinating such a study may be challenging. Given the cost and labour intensive nature of co-heating tests, this would likely require the aggregation of separate co-heating test efforts and the association of tested dwellings with smart meter data.

It may be possible in the future to use smart meter data access portals to obtain data for dwellings which have had co-heating tests performed in the past.

9.3.4 Determining occupant independence of HLC

Further work to confirm that the HLC is independent of occupants is also recommended, as due to limitations of available datasets this could only by inferred from the goodness of fit of the purely physical building model. Establishing beyond doubt that the HLC is independent of building occupants would be a powerful result. Findings in this research have so far supported the theory that the purely physical model on which Deconstruct is based is a reasonable approximation of real energy demand. However it was not possible to prove conclusively that Deconstruct parameters, notably the HLC value, was entirely independent of occupant behaviour. Developing a suitable experimental protocol to verify this would be a valuable avenue for future research. One possible approach would be to compare HLC values for a single building when the occupants have changed. This could be feasible for example with rented accommodation (such as through a council or housing association) where information of tenancy changes could be associated with dwelling energy data and records of repairs and retrofits to the dwelling. If the HLC could be demonstrated to remain approximately constant over time in these circumstances, this would support the theory that HLC is occupant-independent.

9.4 Concluding statement

The Deconstruct method presented in this research demonstrates cost-effective performance analysis on a large scale. In addition to demonstrating good performance on real data, the approach is flexible and open to the incorporation of new data streams generated by the ever increasing diversity of smart devices, appliances, and sensors. As such, it should offer a sound base on which to build further research and analyse new datasets. As an open model, with accessible source code, it should encourage further experimentation and hopefully prove a useful tool in the energy analysis tool kit.

The ability to measure dwelling thermal performance on a large scale in a cost effective manner could fundamentally change the policy and individual outlook on home energy use. Occupants and building managers could have reliable, up-to-date information about their homes, which would help inform retrofit decisions and help provide quality assurance for retrofit works. Policy makers would be able to track the real impact of policy decisions on the change in energy efficiency across the building stock. These factors would help reveal the inherent value-for-money of energy efficiency investments - undoubtedly the simplest and most cost effective approach to reducing greenhouse gas emissions and supporting a sustainable energy system, and also one with the most untapped

potential.

Bibliography

Albert, Adrian, and Ram Rajagopal. 2013. "You Are How You Consume: Mining Structure in Smart Meter Data." Stanford.edu, 1–11.

Allibe, Benoit. 2012. "Modélisation des consommations d'énergie du secteur résidentiel français à long terme." Phd, L'Ecole des Hautes Etudes en Sciences Sociales Spécialité.

Anderson, Brian (BRE), P F Chapman, N G Cutland, C M Dickson, S M Doran, P J Iles, and L D Shorrock. 1997. BREDEM-8. Building Research Establishment.

Andrews, Mike. 2017. "Progress of the CHEESE Project." Bristol: CHEESE Project.

Andrews, Robert W., Joshua S. Stein, Clifford Hansen, and Daniel Riley. 2014. "Introduction to the open source PV LIB for python Photovoltaic system modelling package." In 2014 Ieee 40th Photovoltaic Specialist Conference (Pvsc), 0170–74. IEEE. doi:10.1109/PVSC.2014.6925501.

ArcGIS. 2016. "Shapefiles—ArcGIS Online Help | ArcGIS." https://doc.arcgis.com/en/arcgis-online/reference/shapefiles.htm.

ASHRAE. 2002. "ASHRAE savings calculation." In ASHRAE Guideline 14. Vol. 1986. Ashrae 1990.

ASHRAE Handbook. 2005. "Fundamentals." American Society of Heating, Refrigerating and Air Conditioning Engineers, Atlanta 32: TC 4.7, EnergyCalculations.

Bacher, Peder. 2012. "Models for solar efficient integration of energy."

Bacher, Peder, and Philip Delff Andersen. 2014. "IEA Common Exercise 4: ARX, ARMAX and grey-box models for thermal performance characterization of the test box."

Bacher, Peder, and Henrik Madsen. 2011. "Identifying suitable models for the heat dynamics of buildings." Energy and Buildings 43 (7): 1511–22. doi:10.1016/j.enbuild.2011.02.005.

———. 2013. "Modelling and evaluation of building thermal performance from smart meter

readings," 1–22.

Baker, P.H., and H.A.L. van Dijk. 2008. "PASLINK and dynamic outdoor testing of building components." Building and Environment 43 (2): 143–51. doi:10.1016/j.buildenv.2006.10.009.

Baker, Paul. 2004. "IQ-test—improving quality in testing and evaluation of solar and thermal characteristics of building components." Energy and Buildings 36 (5): 435–41. doi:10.1016/j.enbuild.2004.01.046.

Barnett, Vic., and Toby Lewis. 1978. Outliers in statistical data. Chichester ;;New York: Wiley.

Bauwens, Geert, and Staf Roels. 2014. "Co-heating test: A state-of-the-art." Energy and Buildings 82 (October): 163–72. doi:10.1016/j.enbuild.2014.04.039.

Beard, Chris. 2015. Smart Metering Implementation Programme (SMIP) for Dummies. Chichester: John Wiley & Sons.

BEIS. 2016a. "Energy Consumption in the UK (2016) User Guide." July. London: Department for Business, Energy,; Industrial Strategy.

———. 2016b. "Smart meter roll-out (GB): cost-benefit analysis." https://www.gov.uk/government/publications/smart-meter-roll-out-gb-cost-benefit-analysis.

———. 2016c. "Smart Metering-based Innovation and Building Performance: a BEIS / TEDDINET Exploratory Workshop." https://teddinet.org/collaborating-with-stakeholders-activity.php?s=smart-metering-based-innovation-and-building-performance-a-beis-teddinet-exploratory-workshop.

Bell, Chris. 2016. "Doogal Geocoding." https://www.doogal.co.uk/BatchGeocoding.php.

Berry, Stephen, David Whaley, Kathryn Davidson, and Wasim Saman. 2014. "Near zero energy homes – What do users think?" Energy Policy 73 (October): 127–37. doi:10.1016/j.enpol.2014.05.011.

Biddulph, Phillip, Virginia Gori, Clifford A. Elwell, Cameron Scott, Caroline Rye, Robert Lowe, and Tadj Oreszczyn. 2014. "Inferring the thermal resistance and effective thermal mass of a wall using frequent temperature and heat flux measurements." Energy and Buildings 78 (August): 10–16. doi:10.1016/j.enbuild.2014.04.004.

Birchall, Sarah, Colin Pearson, and Reginald Brown. 2011. "Solid wall insulation Field Trials." Energy Saving Trust. http://www.energysavingtrust.org.uk/Insulation/Solid-wall-insulation.

Birt, Benjamin J., Guy R. Newsham, Ian Beausoleil-Morrison, Marianne M. Armstrong, Neil Saldanha, and Ian H. Rowlands. 2012. "Disaggregating categories of electrical energy end-use from whole-house hourly data." Energy and Buildings 50 (July): 93–102.

doi:10.1016/j.enbuild.2012.03.025.

Boardman, Brenda, C Bank, and G Britain. 2007. “Home Truths: a low-carbon strategy to reduce UK housing emissions by 80% by 2050.” University of Oxford.

Bordass, B, R Cohen, and J Field. 2004. “Energy performance of non-domestic buildings: closing the credibility gap.” Building Performance Congress.

Braun, James, and Nitin Chaturvedi. 2002. “An Inverse Gray-Box Model for Transient Building Load Prediction.” HVAC&R Research 8 (1). Taylor & Francis: 73–99. doi:10.1080/10789669.2002.10391290.

BRE. 1997. “Building Research Establishment Domestic Energy Model (BREDEM).” London: Energy Saving Trust.

———. 2011. SAP 2009 The Government’s Standard Assessment Procedure for Energy Rating of Dwellings. March. Watford: BRE, Garston, Watford WD25 9XX.

———. 2014. SAP 2012 The Government ’ s Standard Assessment Procedure for Energy Rating of Dwellings. October 2013. BRE, Garston, Watford WD25 9XX.

Chambers, Jonathan, and Andrew Stone. 2016. “EPC toolkit.” RCUK-CEE. https://github.com/RCUK-CEE/epctk.

Chambers, Jonathan, Virginia Gori, Phillip Biddulph, Ian Hamilton, Tadj Oreszczyn, and Cliff Elwell. 2015. “How solid is our knowledge of solid walls? - Comparing energy savings through three different methods.” In CISBAT 2015 - International Conference Future Buildings & Districts Sustainability from Nano to Urban Scale.

Chambers, Jonathan, Tadj Oreszczyn, and David Shipworth. 2015. “Quantifying Uncertainty In Grey-box Building Models Arising From Smart Meter Data Sampling Frequency.” Building Simulation Conference 0 (1): 2947–54.

Chapman, J, R Lowe, and R Everett. 1985. “The Pennyland Project.” Milton Keynes: Open University Energy Research Group. http://oro.open.ac.uk/19860/.

Coakley, Daniel, Paul Raftery, and Marcus Keane. 2014. “A review of methods to match building energy simulation models to measured data.” Renewable and Sustainable Energy Reviews 37 (September): 123–41. doi:10.1016/j.rser.2014.05.007.

Cooper, Ian, and Jason Palmer. 2013. “United Kingdom housing energy fact file.”

Crawley, Drury B., Linda K. Lawrie, Frederick C. Winkelmann, W.F. Buhl, Y.Joe Huang, Curtis O. Pedersen, Richard K. Strand, et al. 2001. “EnergyPlus: creating a new-generation building energy

simulation program." Energy and Buildings 33 (4): 319–31. doi:10.1016/S0378-7788(00)00114-6.

Davey, James. 2016. "Cambridge Housing Model and user guide - Publications - GOV.UK." https://www.gov.uk/government/statistics/cambridge-housing-model-and-user-guide.

DCC. 2017. "Data Communications Company." Accessed April 24. https://www.smartdcc.co.uk/about-dcc/.

DCLG. 2013. English Housing Survey Headline report. Department for Communities & Local Government.

———. 2016. "English Housing Survey." Department for Communities; Local Government (DCLG). https://www.gov.uk/government/collections/english-housing-survey.

———. 2017. "Energy Performance of Buildings Data England and Wales." https://epc.opendatacommunities.org/.

DECC. 2009. The UK Low Carbon Transition Plan. London: HM Government.

———. 2011. "Smart Metering Implementation Programme: Industry's Draft Technical Specifications." November. Department of Energy; Climate Change.

———. 2012. "Smart Metering Data Access and Privacy." December. London: DECC.

———. 2013a. "Domestic energy consumption in the UK between 1970 and 2012." In Energy Consumption in the Uk, 1–10. London: DECC.

———. 2013b. "Energy Follow-up Survey 2011 Summary." 289605. London: DECC.

Dempster, A.P., N.M. Laird, and D.B. Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. Vol. 39. 1.

Department for Business Energy & Industrial Strategy. 2016. "Energy Consumption in the UK." Crown Copyright, no. July. www.gov.uk/beis.

Department for Communities and Local Government. 2017. "Energy Performance of Buildings Certificates in England and Wales." http://opendata.cambridgeshireinsight.org.uk/dataset/energy-performance-buildings-certificates-england-and-wales.

Department of Energy and Climate Change. 2014. "Green Deal Assessment Mystery Shopping Research," no. December.

Depuru, Soma Shekara Sreenadh Reddy, Lingfeng Wang, and Vijay Devabhaktuni. 2011. "Smart meters for power grid: Challenges, issues, advantages and status." Renewable and Sustainable Energy Reviews 15 (6): 2736–42. doi:10.1016/j.rser.2011.02.039.

Deurinck, Mieke, Dirk Saelens, and Staf Roels. 2012. "Assessment of the physical part

of the temperature takeback for residential retrofits." Energy and Buildings 52: 112–21. doi:10.1016/j.enbuild.2012.05.024.

Dickson, C.M., J.E. Dunster, S.Z. Lafferty, and L.D. Shorrock. 1996. "BREDEM: Testing monthly and seasonal versions against measurements and against detailed simulation models." Building Services Engineering Research and Technology 17 (3): 135–40. doi:10.1177/014362449601700306.

Doran, Sean, and Bernard Carr. 2008. "Thermal transmittance of walls of dwellings before and after application of cavity wall insulation." 222077. Energy Saving Trust.

Dupret, Muriel, Samuel Martin, Julien Spilemont, Enertech-Presented By, and Muriel Dupret. n.d. "Major role of households specific electricity in total energy consumption of apartment buildings" 2012 (figure 1): 1–12.

D'Oca, Simona, and Tianzhen Hong. 2014. "A data-mining approach to discover patterns of window opening and closing behavior in offices." Building and Environment 82. Elsevier Ltd: 726–39. doi:10.1016/j.buildenv.2014.10.021.

D'Oca, Simona, Valentina Fabi, Stefano P. Corgnati, and Rune Korsholm Andersen. 2014. "Effect of thermostat and window opening occupant behavior models on energy use in homes." Building Simulation 7 (6): 683–94. doi:10.1007/s12273-014-0191-6.

EC. 2010. "Interpretative Note on Directive 2009/72/Ec Concerning Common Rules for the Internal Market in Electricity and Directive 2009/73/Ec Concerning Common Rules for the Internal Market in Natural Gas." European Commisson.

———. 2014. "Over 70% European consumers to have a smart meter for electricity by 2020." https://ec.europa.eu/jrc/en/news/over-70-percent-european-consumers-have-smart-meter-electricity-2020.

Edwards, Richard E., Joshua New, and Lynne E. Parker. 2012. "Predicting future hourly residential electrical consumption: A machine learning case study." Energy and Buildings 49: 591–603. http://www.sciencedirect.com/science/article/pii/S0378778812001582.

Elwell, Clifford A., Phillip Biddulph, Robert Lowe, and Tadj Oreszczyn. 2015. "Determining the impact of regulatory policy on UK gas use using Bayesian analysis on publicly available data." Energy Policy 86 (January 1990). Elsevier: 770–83. doi:10.1016/j.enpol.2015.08.020.

En, B S. 2006. "Gas-fired central heating boilers — Type B boilers of nominal heat input exceeding 70 kW but not exceeding 300 kW." Engineering 3 (1).

Energy Saving Trust. 2008. "Measurement of Domestic Hot Water Consumption in Dwellings."

Energy Savings Trust.

EPFL. 2013. “CISBAT 2013.” In. Lausanne.

EST. 2017. “Domestic Energy Assessor Training Courses (DEA) | Energy Trust.” https://www.energy-trust.co.uk/collections/domestic-energy-assessor-dea-training.

ETI. 2016. “Housing Retrofits.” Energy Technologies Institute.

EU. 2006. “Energy end-use efficiency and energy services.” Brussels: European Commission. http://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:32006L0032.

———. 2007. “The 2020 climate and energy package - European Commission.” http://ec.europa.eu/clima/policies/package/i

———. 2010. “Directive 2010/31/EU of the European Parliament and of the Council of 19 May 2010 on the energy performance of buildings.” Official Journal of the European Union, 13–35. doi:doi:10.3000/17252555.L_2010.153.eng.

———. 2012. “Energy Efficiency Directive 2012.” Brussels: European Commission. http://eur-lex.europa.eu/legal-content/en/TXT/?uri=celex:32012L0027.

———. 2014. “2030 framework for climate and energy policies - European Commission.” http://ec.europa.eu/clima/policies/2030/index{_}en.htm.

European Parliament. 2004. “Directive 2004/22/EC of the European Parliament and of the Council of 31 March 2004.” Official Journal of the European Union 2003 (807): 1–80.

Eurostat. 2016. “NUTS - Nomenclature of territorial units for statistics.” http://ec.europa.eu/eurostat/web/nuts/overview.

Everett, R., A. Horton, and J. Doggart. 1985a. “Linford Low Energy Houses Project Summary.” http://oro.open.ac.uk/19861/.

Everett, R., a. Horton, and J. Doggart. 1985b. “Linford Low Energy Houses.” http://oro.open.ac.uk/19861/.

Everett, Robert. 1985. “Rapid Thermal Calibration of Houses.” http://www.aivc.org/resource/rapid-thermal-calibration-houses.

Experimental Office of National Statistics. 2016. “Smart Meters Quarterly Report to end September 2016 Great Britain,” no. September. www.gov.uk/government/publications.

Fanger, P. O. 1970. Thermal comfort: analysis and applications in environmental engineering. New York : McGraw-Hill.

Fels, Margaret. 1986. “PRISM: An introduction.” Energy and Buildings 9 (1-2): 5–18.

doi:10.1016/0378-7788(86)90003-4.

Fels, Margaret F., Miriam L. Goldberg, and Michael L. Lavine. 1986. “Exploratory scorekeeping for oil-heated houses.” Energy and Buildings 9 (1-2): 127–36. doi:10.1016/0378-7788(86)90014-9.

Fox, Matthew, David Coley, Steve Goodhew, and Pieter de Wilde. 2014. “Thermography methodologies for detecting energy related building defects.” Renewable and Sustainable Energy Reviews 40 (December): 296–310. doi:10.1016/j.rser.2014.07.188.

FrictionlessData. 2017. “Data Packages.” http://frictionlessdata.io/data-packages/.

Friedman, Jerome H. 1991. “Multivariate Adaptive Regression Splines.” The Annals of Statistics 19 (1). Institute of Mathematical Statistics: 1–67. http://projecteuclid.org/euclid.aos/1176347963.

Gilbertson, Jan, and Geoff Green. 2008. Warm Front Better Health. DEFRA.

Goldberg, ML, and MF Fels. 1986. “Refraction of PRISM Results into Components of Saved Energy.” Energy and Buildings 9: 169–80. http://www.sciencedirect.com/science/article/pii/0378778886900186.

Gori, Virginia. 2016. “Personal Correspondence.”

Greening, Lorna, David L. Greene, and Carmen Difiglio. 2000. “Energy efficiency and consumption — the rebound effect — a survey.” Energy Policy 28 (6): 389–401. http://www.sciencedirect.com/science/article/pii/S0301421500000215.

Gregg, Matt. 2014. “A quiet revolution : Mapping energy use in low carbon communities,” no. December: 1–8.

Hamilton, Ian G., David Shipworth, Alex J. Summerfield, Philip Steadman, Tadj Oreszczyn, and Robert Lowe. 2014. “Uptake of energy efficiency interventions in English dwellings.” Building Research & Information 42 (3). Taylor & Francis: 255–75. doi:10.1080/09613218.2014.867643.

Hamilton, Ian G., Alex J. Summerfield, David Shipworth, J. Philip Steadman, Tadj Oreszczyn, and Robert J. Lowe. 2016a. “Energy efficiency uptake and energy savings in English houses: A cohort study.” Energy and Buildings 118 (April). Elsevier Ltd: 259–76.

Hamilton, Ian, Aidan O’Sullivan, Gesche Huebner, David Shipworth, Alex Summerfield, Tadj Oreszczyn, and M Davies. 2016b. “Old and cold: findings on the determinants of indoor temperatures in English dwellings during cold conditions.”

HMGovernment. 2008. “Climate Change Act 2008.” Statute Law Database. http://www.

legislation.gov.uk/ukpga/2008/27/contents.

———. 2010. “2050 Pathways Analysis.” HM Government.

Hoffmann, Caroline, and Achim Geissler. 2017. “Dem Prebound Effekt auf der Spur – Differenzen zwischen dem Heizwärmeverbrauch und dem rechnerisch ermittelten Heizwärmebedarf bei Bestandsgebäuden (Wohnen).” Bauphysik 39 (3): 159–74. doi:10.1002/bapi.201710022.

Hogan, Robin. 2006. “How to combine errors,” no. June: 1–2.

Hoyer, S, and J Hamman. 2016. “xarray: N-D labeled arrays and datasets in Python.” In Prep, J. Open Res. Software.

Huebner, Gesche M., Ian Hamilton, Zaid Chalabi, David Shipworth, and Tadj Oreszczyn. 2015. “Explaining domestic energy consumption – The comparative contribution of building factors, socio-demographics, behaviours and attitudes.” Applied Energy 159. Elsevier Ltd: 589–600. doi:10.1016/j.apenergy.2015.09.028.

Huebner, Gesche M., Megan McMichael, David Shipworth, Michelle Shipworth, Mathieu Durand-Daubin, and Alex Summerfield. 2013. “The reality of English living rooms – A comparison of internal temperatures against common model assumptions.” Energy and Buildings 66 (November). Elsevier B.V.: 688–96. doi:10.1016/j.enbuild.2013.07.025.

Huebner, Gesche M., Megan McMichael, David Shipworth, Michelle Shipworth, Mathieu Durand-Daubin, and Alex J. Summerfield. 2014. “The shape of warmth: temperature profiles in living rooms.” Building Research & Information, June. Routledge, 1–12. doi:10.1080/09613218.2014.922339.

Hughes, Martin, Jason Palmer, Vicky Cheng, and David Shipworth. 2013. “Sensitivity and uncertainty analysis of England’s housing energy model.” Building Research & Information 41 (2). Routledge: 156–67. doi:10.1080/09613218.2013.769146.

———. 2014. “Global sensitivity analysis of England’s housing energy model.” Journal of Building Performance Simulation, June. Taylor & Francis, 1–12. doi:10.1080/19401493.2014.925505.

Hulme, Jack, Adele Beaumont, and Claire Summers. 2013a. “Energy Follow-up Survey 2011 Report 4 : Main heating systems.” 286733.

———. 2013b. “Energy Follow-up Survey 2011, Report 11 : Methodology,” no. 288851.

———. 2013c. “Energy Follow-Up Survey 2011, Report 5 : Secondary heating systems.” BRE on Behalf of DECC, no. 286733.

Hunter, John D. 2007. “Matplotlib: A 2d Graphics Environment.” Computing in Science &

Engineering 9 (3): 90–95. doi:10.1109/MCSE.2007.55.

IEA. 2010. World Energy Outlook 2010. World Energy Outlook. Paris, France: OECD Publishing. doi:10.1787/weo-2010-en.

———. 2013. Energy efficiency market report 2013 : market trends and medium-term prospects. Paris, France: OECD/IEA.

IET. 2013. "Smart Meter Roll-Out." 211014. London: Instituteion of Engineering; Technology.

Iglewicz, Boris, and David Caster Hoaglin. 1993. How to Detect and Handle Outliers. ASQC Quality Press.

IHE. 2014. "Fuel Poverty and Cold Home-related Health Problems," no. September.

Ineichen, Pierre. 2008. "Comparison and validation of three global-to-beam irradiance models against ground measurements." Solar Energy 82 (6): 501–12. doi:10.1016/j.solener.2007.12.006.

IPCC. 2013. "Fifth Assessment Report - Climate Change 2013." http://www.ipcc.ch/report/ar5/wg1/.

ISO. 2010. Energy performance of buildings - Calculation of energy use for space heating and cooling (ISO 13790:2008) Performance. 5972936.

Jack, Richard. 2015. "Building Diagnostics : Practical Measurement of the Fabric Thermal Performance of Houses," no. July: 254. doi:10.13140/RG.2.1.1731.4641.

Janssens, Arnold. 2014. Reliable building energy performance characterisation based on full scale dynamic measurements in Buildings Background : Renewed interest in full scale testing Interest. May.

Jiménez, M. J., B. Porcar, and M. R. Heras. 2009. "Application of different dynamic analysis approaches to the estimation of the building component U value." Building and Environment 44 (2): 361–67. doi:10.1016/j.buildenv.2008.03.010.

Joint Committee for Guides in Metrology (JCGM). 2008. "JCGM 101:2008 Evaluation of measurement data – Supplement 1 to the ' Guide to the expression of uncertainty in measurement ' – Propagation of distributions using a Monte Carlo method."

Jones, Eric, Travis Oliphant, Pearu Peterson, and Others. 2015. "SciPy: Open source scientific tools for Python." http://www.scipy.org/.

Juhl, Rune, Niels Rode Kristensen, Peder Bacher, Jan Kloppenbord, and Henrik Madsen. 2013. "Grey-box modeling of the heat dynamics of a building with CTSM-R."

Kelly, Scott, Doug Crawford-Brown, and Michael G. Pollitt. 2012. "Building performance eval-

uation and certification in the UK: Is SAP fit for purpose?" Renewable and Sustainable Energy Reviews 16 (9): 6861–78. http://www.sciencedirect.com/science/article/pii/S1364032112004595.

Kim, Jin-Ho, and Anastasia Shcherbakova. 2011. "Common failures of demand response." Energy 36 (2): 873–80. doi:10.1016/j.energy.2010.12.027.

Lander, Dave F. 2012. "Accuracy of CV Determination Systems for Calculation of FWACV." DLC/0030. Tamworth: EMIB.

Laurent, Marie Hélène. 2014. "Personal Correspondence." Fontainebleau: EDF R&D.

Li, Francis G. N., A.Z.P. Smith, Phillip Biddulph, Ian G. Hamilton, Robert Lowe, Anna Mavrogianni, Eleni Oikonomou, et al. 2014. "Solid-wall U -values: heat flux measurements compared with standard assumptions." Building Research & Information, October. Routledge, 1–15. doi:10.1080/09613218.2014.967977.

Loutzenhiser, P.G., H. Manz, C. Felsmann, P.A. Strachan, T. Frank, and G.M. Maxwell. 2007. "Empirical validation of models to compute solar irradiance on inclined surfaces for building energy simulation." Solar Energy 81 (2): 254–67. doi:10.1016/j.solener.2006.03.009.

Love, Jennifer. 2014. "Understanding the Interactions Between Occupants, Heating Systems, and Building Fabric in the Context of Social Housing." PhD thesis.

Lowe, R.J., and C.J. Gibbons. 1988. "Passive solar houses: Availability of weather suitable for calibration in the UK." Building Services Engineering Research and Technology 9 (3). Sage PublicationsSage CA: Thousand Oaks, CA: 127–32. doi:10.1177/014362448800900306.

Madsen, Henrik. 2015. "Thermal Performance Characterization using Time Series Data ; IEA EBC Annex 58 Guidelines," no. January 2016. doi:10.13140/RG.2.1.1564.4241.

Majcen, D., L. C M Itard, and H. Visscher. 2013. "Theoretical vs. actual energy consumption of labelled dwellings in the Netherlands: Discrepancies and policy implications." Energy Policy 54 (March): 125–36. doi:10.1016/j.enpol.2012.11.008.

McKinney, Wes. 2010. "Data Structures for Statistical Computing in Python."

Met Office. 2016. "Cartopy: a cartographic python library with a matplotlib interface." Exeter, Devon. http://scitools.org.uk/cartopy.

MetOffice. 2013. "Met Office Integrated Data Archive System (MIDAS)." http://catalogue.ceda.ac.uk/uuid/220a65615218d5c9cc9e4785a3234bd0.

———. 2015. "Personal Correspondence."

———. 2016. "Met Office Numerical Weather Prediction models." Met Office, FitzRoy Road,

Exeter, Devon, EX1 3PB, United Kingdom.

Meyer, Leo. 1981. “Energiebesparing in de Sociale Woningbouw.” PhD thesis, University of Groeningen.

Moon, Nick, Polly Hollings, and Gfk Nop. 2014. “Complaints to Energy Companies: Exploring How Energy Companies Handle Customer Complaints.” London: Ofgem.

Nascimento, Ricardo Mion do, Ana Paula Oening, Debora Cintia Marcilio, Alexandre Rasi Aoki, Eloy de Paula Rocha, and Jefferson Marcelo Schiochet. 2012. “Outliers’ detection and filling algorithms for smart metering centers.” In PES T&D 2012, 1–6. IEEE. doi:10.1109/TDC.2012.6281659.

NationalGrid. 2015. “Calorific Value Information | National Grid.” http://www2.nationalgrid.com/UK/Industry-information/Gas-transmission-operational-data/calorific-value-description/.

Navigant Research. 2016. “Market Data: Smart Meters.” http://www.navigantresearch.com/research/market-data-smart-meters.

Ng, Andrew. 2013. “Machine Learning | Coursera.” https://www.coursera.org/course/ml.

NIST. 2013. “Detection of Outliers.” http://www.itl.nist.gov/div898/handbook/eda/section3/eda35h.htm.

Ofgem. 2000. “November 2000 Gas energy measurement A consultation document.” November. London.

———. 2013. “Transition to smart meters | Ofgem.” https://www.ofgem.gov.uk/electricity/retail-market/metering/transition-smart-meters.

ONS. 2016. “Code-Point Open - locates every postcode unit in GB with precision.” https://www.ordnancesurvey.co.uk/business-and-government/products/code-point-open.html.

OpenDoorLogistics. 2016. “Open Door Logistics postcode boundaries.” http://www.opendoorlogistics.com/data/.

Ordnance Survey. 2016. “Ordnance Survey: Britain’s mapping agency.” https://www.ordnancesurvey.co.uk/.

Orehounig, Kristina, Sokol Dervishi, and Ardeshir Mahdavi. 2011. “Comparison of Computed and Measured Irradiance on Building Surfaces.” In IBPSA, 14–16. 2010.

Oreszczyn, Tadj. 2014. “Personal Correspondence.” London: UCL Energy Institute.

Oreszczyn, Tadj, and Robert Lowe. 2010. “Challenges for energy and buildings research: objectives, methods and funding mechanisms.” Building Research & Information 38 (1). Routledge:

107–22. doi:10.1080/09613210903265432.

Oreszczyn, Tadj, Sung H. Hong, Ian Ridley, and Paul Wilkinson. 2006. “Determinants of winter indoor temperatures in low income households in England.” Energy and Buildings 38 (3): 245–52. http://www.sciencedirect.com/science/article/pii/S0378778805000964.

Orr, Georgina, Tom Lelyveld, and Simon Burton. 2009. “Final Report : In-situ monitoring of efficiencies of condensing boilers and use of secondary heating,” no. June: 44.

Palmer, Jason, Amy Tillson, and Peter Armitage. 2013. “Comparing the Cambridge Housing Model against the National Energy Efficiency Data-Framework and Meter Readings,” no. July.

Parliament. 2008. “Climate Change Act 2008.” London: UK. http://www.publications.parliament.uk/pa/cm200708/cmbills/129/08129.i-v.html.

Pedregosa, Fabian, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, et al. 2011. “Scikit-learn: Machine Learning in Python.” Journal of Machine Learning Research 12 (Oct): 2825–30.

Perez, Fernando, and Brian E. Granger. 2007. “IPython: A System for Interactive Scientific Computing.” Computing in Science & Engineering 9 (3): 21–29. doi:10.1109/MCSE.2007.53.

Python Software Foundation. 2015. “Python 3.5.” Python Software Foundation. www.python.org.

Rabl, A. 1988. “Parameter Estimation in Buildings: Methods for Dynamic Analysis of Measured Energy Use.” Journal of Solar Energy Engineering 110: 52–66.

Raw, Gary, David Ross, and Andrew Cripps. 2011. “Energy Demand Research Project : Final Analysis.” June. St Albans: AECOM.

Reda, Ibrahim, and Afshin Andreas. 2004. “Solar position algorithm for solar radiation applications.” Solar Energy 76 (5): 577–89. doi:10.1016/j.solener.2003.12.003.

Reno, Matthew J, Clifford W Hansen, and Joshua S Stein. 2012. “Global Horizontal Irradiance Clear Sky Models: Implementation and Analysis.” SANDIA REPORT Sand2012-2389 Unlimited Release Printed March 2012, no. March: 1–66. doi:10.2172/1039404.

Roberts, Simon. 2008. “Altering existing buildings in the UK.” Energy Policy 36 (12): 4482–6. doi:10.1016/j.enpol.2008.09.023.

Roels, Staf. 2012. “Annex 58 - Reliable Building Energy Performance Characterisation Based on Full Scale Dynamic Measurements.” The International Energy Agency Annex 58 (June): 22.

———. 2017. “EBC Annex 71 Building Energy Performance Assessment Based on In-situ Mea-

surements." http://www.iea-ebc.org/projects/ongoing-projects/ebc-annex-71/.

Rosenow, Jan, and Nick Eyre. 2016. "A post mortem of the Green Deal: Austerity, energy efficiency, and failure in British energy policy." Energy Research & Social Science 21 (November): 141–44. doi:10.1016/j.erss.2016.07.005.

Rosenow, Jan, and Ray Galvin. 2013. "Evaluating the evaluations: Evidence from energy efficiency programmes in Germany and the UK." Energy and Buildings 62: 450–58. http://www.sciencedirect.com/science/article/pii/S0378778813001898.

Ruch, David K., and David E. Claridge. 1993. "A development and comparison of NAC estimates for linear and change-point energy models for commercial buildings." Energy and Buildings 20 (1): 87–95. doi:10.1016/0378-7788(93)90041-R.

Ruch, David, and David E. Claridge. 1992. "A Four-Parameter Change-Point Model for Predicting Energy Consumption in Commercial Buildings." Journal of Solar Energy Engineering 114 (2). American Society of Mechanical Engineers: 77. doi:10.1115/1.2929993.

Saha, Suranjana, Shrinivas Moorthi, Hua-Lu Pan, Xingren Wu, Jiande Wang, Sudhir Nadiga, Patrick Tripp, et al. 2010. "The NCEP Climate Forecast System Reanalysis." Bulletin of the American Meteorological Society 91 (8). American Meteorological Society: 1015–57. doi:10.1175/2010BAMS3001.1.

Sanders, Chris, and Mark Phillipson. 2006. "Review of Differences between Measured and Theoretical Energy Savings for Insulation Measures." December.

Scottish Government. 2009. Modelling Greenhouse Gas Emissions from Scottish Housing: Final Report. Edinburgh: Scottish Government. http://www.scotland.gov.uk/Publications/2009/10/08143041/3.

Shipworth, Michelle. 2011. "Thermostat settings in English houses: No evidence of change between 1984 and 2007." Building and Environment 46 (3): 635–42. doi:10.1016/j.buildenv.2010.09.009.

Shiret, Alan (BRE), and John Hayton (BRE). 2010. "Laboratory tests on domestic gas and oil boilers." June.

Siviour, JB. 1981. Experimental Thermal Calibration of Houses. Chester: Electricity Council Research Centre. http://www.aivc.org/resource/experimental-thermal-calibration-houses.

Soderstrom, T., A Rabl, G Fracastoro, E. Nino, J.J Bloem, Henrik Madsen, H. Melgaard, et al. 1996. SYSTEM IDENTIFICATION APPLIED TO BUILDING PERFORMANCE DATA. Edited by J. J. Bloem. Luxembourg: Institute for Systems Engineering; Informatics System Identification.

Software Sustainability Institute. 2016. "The Software Sustainability Institute." https://www.

software.ac.uk/.

Stafford, A, M Bell, and C Gorse. 2012. "Building Confidence – A working paper." http://www.lowcarbonfutures.org/sites/default/files/2781{_}Building{_}confidence{_}report{_}lores.pdf.

Stafford, A., C. Gorse, and L. Shao. 2011. "The Retrofit Challenge: Delivering Low Carbon Buildings." Leeds: Centre for Low Carbon Futures.

Stamp, Samuel Francis. 2015. "Assessing Uncertainty in Co-heating Tests : Calibrating a Whole Building Steady State Heat Loss Measurement Method."

Stone, Andrew, Tadj Oreszczyn, David Shipworth, and Phill Biddulph. 2013. "Global Sensitivity Analysis of a BREDEM Based Domestic Energy Model - DRAFT."

Stone, Andrew, David Shipworth, Phill Biddulph, and Tadj Oreszczyn. 2014. "Key factors determining the energy rating of existing English houses." Building Research & Information 42 (6): 725–38. doi:10.1080/09613218.2014.905383.

Summerfield, A.J., T Oreszczyn, J Palmer, I.G. Hamilton, and R.J. Lowe. 2015a. "Comparison of empirical and modelled energy performance across age-bands of three-bedroom dwellings in the UK." Energy and Buildings 109: 328–33. doi:10.1016/j.enbuild.2015.09.050.

Summerfield, A.J., T. Oreszczyn, I.G. Hamilton, D. Shipworth, G.M. Huebner, R.J. Lowe, and P. Ruyssevelt. 2015b. "Empirical variation in 24-h profiles of delivered power for a sample of UK dwellings: Implications for evaluating energy savings." Energy and Buildings 88 (February): 193–202. doi:10.1016/j.enbuild.2014.11.075.

Summerfield, Alex J., R. J. Lowe, and T. Oreszczyn. 2010. "Two models for benchmarking UK domestic delivered energy." Building Research & Information 38 (1): 12–24. doi:10.1080/09613210903399025.

Summerfield, Alex J., R.J. Lowe, H.R. Bruhns, J.A. Caeiro, J.P. Steadman, and T. Oreszczyn. 2007. "Milton Keynes Energy Park revisited: Changes in internal temperatures and energy usage." Energy and Buildings 39 (7): 783–91. http://www.sciencedirect.com/science/article/pii/S0378778807000576.

Summerfield, Alex J., Tadj Oreszczyn, and Sung-hyon Hong. 2013. "UK Dwelling Temperatures DRAFT," 1–13.

Sunikka-Blank, Minna, and Ray Galvin. 2012. "Introducing the prebound effect: the gap between performance and actual energy consumption." Building Research & Information 40 (3). Routledge: 260–73. doi:10.1080/09613218.2012.690952.

Swan, Lukas G., and V. Ismet Ugursal. 2009. "Modeling of end-use energy consumption in the residential sector: A review of modeling techniques." Renewable and Sustainable Energy Reviews

13 (8): 1819–35. http://www.sciencedirect.com/science/article/pii/S1364032108001949.

UCAR. 2016. “Network Common Data Form (NetCDF).” http://www.unidata.ucar.edu/software/netcdf/.

UK. 2012. “Energy Performance Certificate (EPC).” London.

UK Data Service. 2014a. “Energy Demand Research Project: Early Smart Meter Trials, 2007-2010.” AECOM Building Engineering. https://discover.ukdataservice.ac.uk/catalogue/?sn=7591.

———. 2014b. “Smart Meter Energy Demand Research Project : Data Release,” no. 7591: 2007–10.

UK-GBC. 2014. “New Build | UK Green Building Council.” http://www.ukgbc.org/content/new-build.

UKGovernment. 2017. “Buying or selling your home: Energy Performance Certificates - GOV.UK.” Accessed August 8. https://www.gov.uk/buy-sell-your-home/energy-performance-certificates.

USwitch. 2017. “Energy performance certificates - joining the EPC register.” Accessed April 24. https://www.uswitch.com/energy-saving/guides/energy-performance-certificates/{\#}step6.

“Violin Plot.” 2017. Accessed May 21. http://www.datavizcatalogue.com/methods/violin{_}plot.html.

Walker, I. S., and D. J. Wilson. 1998. “Field Validation of Algebraic Equations for Stack and Wind Driven Air Infiltration Calculations.” HVAC&R Research, March. Taylor & Francis Group. http://www.tandfonline.com/doi/abs/10.1080/10789669.1998.10391395.

Walt, Stéfan van der, S Chris Colbert, and Gaël Varoquaux. 2011. “The NumPy Array: A Structure for Efficient Numerical Computation.” Computing in Science & Engineering 13 (2): 22–30. doi:10.1109/MCSE.2011.37.

Which. 2014. “EPC explained.” http://www.which.co.uk/money/mortgages-and-property/guides/how-to-sell-a-house/epc-explained/.

Wilkinson, Paul, Megan Landon, Ben Armstrong, Simon Stevenson, Sam Pattenden, Martin Mckee, and Tony Fletcher. 2001. “Cold comfort: The social and environmental determinants of excess winter deaths in England , 1986-96,” 34. http://www.jrf.org.uk/publications/cold-comfort-social-and-environmental-determinants-excess-winter-deaths-england-1986-19.

Wilkinson, Paul, Kirk R Smith, Sean Beevers, Cathryn Tonne, and Tadj Oreszczyn. 2007. “Energy, energy efficiency, and the built environment.” Lancet 370 (9593): 1175–87. doi:10.1016/S0140-

6736(07)61255-0.

Wingfield, Jez. 2017. “Personal Correspondence.”

Zwillinger, D., and S Kokoska. 2000. CRC Standard Probability and Statistics Tables and Formulae. New York, New York, USA: Chapman & Hall.

Appendix A

Internal temperature dependence on external temperature

Consider a standard dwelling heating pattern consisting of on and off periods. Assume that during the on periods the internal temperature is constant T_{on}. During off periods, the dwelling will cool down following an exponential decay according to Newton's law of cooling. Define a 'limit' temperature T_{sc} which is the theoretical steady-state temperature a dwelling would be at if there was no heating applied. This is the temperature that the decay curve will trend towards. The internal temperature over time during cool-down will therefore start at T_{on} and decay exponentially towards T_{sc} over a period of time governed by the dwelling's time constant τ:

$$T_{in}(t) = T_{sc} + (T_{on} - T_{sc})e^{-t/\tau}$$

The heating profile consists of sets of on and off times, t_{ons} and t_{offs}. Summarise the daily internal temperature profile as a piecewise function:

$$T_{in} = \begin{cases} T_{on} & \rightarrow t \in t_{ons} \\ T_{sc} + (T_{on} - T_{sc})e^{-t/\tau} & \rightarrow t \in t_{offs} \end{cases}$$

The mean 24hr temperature $\overline{T_{in}}$ is (with time measured in hours)

$$\overline{T_{in}} = 1/24 \sum_{j} \Delta t_j \overline{T_j}$$

where for each segment j of the heating profile, Δt_j is the duration of the segment and $\overline{T_j}$ is the mean temperature in that segment. The mean temperature when heating is on is simply T_{on} -

this should ideally be the set-point, but for a whole dwelling average may be somewhat different. Calculate the mean temperature during a given cool off period, defining a time t such that the start of the cool-down period is at 0 and the end of the time period is the time the heating comes on again t_{on}. This can be generalised to arbitrary off and on times, but using these limits simplifies the exercise.

$$\begin{aligned}
\overline{T_{in,cool}} &= \frac{1}{t_{on}} \int_0^{t_{on}} T_{sc} + (T_{on} - T_{sc}) e^{-t/\tau} dt \\
&= \frac{1}{t_{on}} \left[T_{sc} t - \tau (T_{on} - T_{sc}) e^{-t/\tau} \right]_0^{t_{on}} \\
&= \frac{1}{t_{on}} (\left[T_{sc} t_{on} - \tau (T_{on} - T_{sc}) e^{-t_{on}/\tau} \right] - \left[T_{sc} 0 - \tau (T_{on} - T_{sc}) e^{0/\tau} \right]) \\
&= \frac{1}{t_{on}} (\left[T_{sc} t_{on} - \tau (T_{on} - T_{sc}) e^{-t_{on}/\tau} \right] + \tau (T_{on} - T_{sc})) \\
&= T_{sc} + \frac{\tau}{t_{on}} (T_{on} - T_{sc}) \left(1 - e^{-t_{on}/\tau} \right)
\end{aligned}$$

The mean temperature in the cooldown period is the limit temperature T_{sc} plus a factor that depends on the difference between the heating and the limit temperature and the time passed relative to the building's time constant.

The mean temperature over one day is then the sum of the means for different parts of the cycle.

$$\overline{T_{in}} = \frac{1}{24} (\sum_j \Delta t_j T_{on} + \sum_k \Delta t_k (T_{sc} + \frac{\tau}{t_k} (T_{on} - T_{sc}) (1 - e^{-t_k/\tau}))$$

We are interested in developing a model for T_{in} as a function of T_{ex}. The first step is to determine how T_{in} changes as a function of T_{ex} i.e. the first derivative of the mean.

The terms in the heating periods drop out immediately since they are constant, the derivative of the mean during cooldown period of duration t_k is as follows

$$\begin{aligned}
\frac{d\overline{T_{in,cool}}}{dT_{ex}} &= \frac{d}{dT_{ex}} \left(T_{sc} + \frac{\tau}{t_k} (T_{on} - T_{sc}) (1 - e^{-t_k/\tau} \right) \\
&= \frac{d}{dT_{ex}} T_{sc} + \frac{\tau}{t_k} \frac{d}{dT_{ex}} (T_{on} - T_{sc}) (1 - e^{-t_k/\tau}) \\
&= \frac{d}{dT_{ex}} T_{sc} + \frac{\tau}{t_k} \frac{d}{dT_{ex}} T_{on} (1 - e^{-t_k/\tau}) - \frac{\tau}{t_k} \frac{d}{dT_{ex}} T_{sc} (1 - e^{-t_k/\tau})
\end{aligned}$$

We assume that the heating system maintains a constant temperature during the heating periods, i.e. that $\frac{dT_{on}}{dT_{ex}} = 0$

$$\frac{d\overline{T_{in,cool}}}{dT_{ex}} = \frac{dT_{sc}}{dT_{ex}} - \frac{\tau}{t_k} \frac{dT_{sc}}{dT_{ex}} (1 - e^{-t_k/\tau})$$

To summarise:

$$\frac{d\overline{T_{in,cool}}}{dT_{ex}} = \frac{dT_{sc}}{dT_{ex}}(1 - \frac{\tau}{t_k} + \frac{\tau}{t_k}e^{-t_k/\tau})$$

For a given heating pattern, for each cooldown period k, the term $(1 - \frac{\tau}{t_k} + \frac{\tau}{t_k}e^{-t_k/\tau})$ is constant since the duration of the cooldown period is set by the heating pattern and the time constant is a property of the building.

The derivative with respect to external temperature for a daily heating profile is then

$$\frac{d\overline{T_{in}}}{dT_{ex}} = \frac{1}{24}(\sum_k \Delta t_k(\frac{dT_{sc}}{dT_{ex}}(1 - \frac{\tau}{t_k} + \frac{\tau}{t_k}e^{-t_k/\tau}))$$

The key term is $\frac{dT_{sc}}{dT_{ex}}$. A simple first order approximation for T_{sc} can be obtained using a steady-state heat balance with linearised losses:

$$\begin{aligned} \phi_{gain} - \phi_{loss} &= 0 \\ H_{tr}(T_{sc} - T_{ex}) &= \phi_{gain} \\ T_{sc} &= T_{ex} + \phi_{gain}/H_{tr} \end{aligned}$$

The result being that the steady-state mean internal temperature without heating is the external temperature plus a factor that depends on other incidental gains (solar gains, appliances) and the heat loss coefficient. With this approximation, if we further assume that incidental gains do not depend on T_{ex} then $\frac{dT_{sc}}{dT_{ex}} = 1$ and $\frac{d\overline{T_{in}}}{dT_{ex}}$ is

$$\frac{d\overline{T_{in}}}{dT_{ex}} = \frac{1}{24}\left(\sum_k \Delta t_k(1 - \frac{\tau}{t_k} + \frac{\tau}{t_k}e^{-t_k/\tau})\right)$$

where all times parameters are in hours. This demonstrates that the covariance of mean daily internal temperature with external temperature ($\frac{d\overline{T_{in}}}{dT_{ex}}$) is a function of the length of the heating periods and the temperature decay rate.

Appendix B

Building physics - base equations

This section outlines a number of building physics equations and derivations relevant to the Deconstruct model theory.

B.1 Solar gains through transparent and opaque surfaces

Solar irradiance is an important contributor to dwelling thermal balance (Orehounig, Dervishi, and Mahdavi 2011). Energy is absorbed by a building both through the heating of solid surfaces and via transparent windows transmit and retain radiation inside the building (Bauwens and Roels 2014).

The net radiative energy $\Phi_{sol,k}$ (kW) absorbed by a surface K with area A_k is a function of the dimensionless absorptivity α_k of the surface and the irradiance incident on the surface from all angles at a given time $I_{k,t,tot}$ (kW). For transparent components like windows the 'g-value' is used to indicate their transmittance of solar energy, on the assumption that admitted irradiance will be transmitted into the building. The total heat gain from an element is (eq. B.1):

$$\Phi_{sol,k} = \alpha_k A_k I_{k,tot} + g_k A_k I_{k,tot} \tag{B.1}$$

The total irradiance $I_{k,tot}$ will change over time as the sun moves through the sky and the pattern of diffuse irradiance from clouds and from the surrounding environment changes (Loutzenhiser et al. 2007). The total irradiance over the day for a surface will be a function of the orientation of the surface, the weather conditions, and time.

In practice, the fraction of the energy absorbed by an opaque surface and transmitted inside the building is small compared to the energy absorbed through transparent elements. We will include

this fraction for theoretical completeness, by defining parameter $a_k = f\alpha_k + g_k$, that is, the sum of the fraction f of surface absorbed energy transmitted into the building plus the fraction of energy transmitted if the element is transparent. The solar gains over all building elements as a function of time $\Phi_{sol}(t)$ is therefore:

$$\Phi_{sol}(t) = \sum a_k A_k I_{k,tot} \tag{B.2}$$

Taking the integral over time of the sum of exposed surfaces, the energy absorbed is a function of the number and orientation of the sun-exposed surfaces of a dwelling and the total (direct plus diffuse) irradiance reaching those surfaces over time:

$$\begin{aligned} \Phi_{sol} &= \frac{1}{t_2 - t_1} \int_{t_1}^{t_2} \Phi_{sol}(t) dt \\ &= \frac{1}{t_2 - t_1} \int_{t_1}^{t_2} \sum a_k A_k I_{k,tot}(t) dt \\ &= \sum a_k A_k \frac{1}{t_2 - t_1} \int_{t_1}^{t_2} I_{k,tot}(t) dt \\ &= \sum a_k A_k \overline{I_{k,tot}} \end{aligned} \tag{B.3}$$

where $\overline{I_{k,tot}}$ is the mean daily irradiance on surface i. Bauwens and Roels (2014) describe the simplification eq. B.4 used in co-heating tests.

$$\Phi_{sol} = A_{sol} I_{sol} \tag{B.4}$$

where A_{sol} in m2 is the effective solar aperture of the building and I_{sol} is the received solar irradiance (kW/m2). The equivalence between eq. B.3 and eq. B.4 can be thought of as finding an average $\overline{I_{k,tot}}$ over all i elements such that

$$\begin{aligned} \Phi_{sol} &= \sum a_k A_k \overline{I_{k,tot}} \\ &= \overline{I_{sol}} \sum a_k A_k \end{aligned} \tag{B.5}$$

The effective solar aperture $A_{sol} = \sum a_k A_k$ does not therefore have a singular physical meaning, but can be understood as an approximation lumping together a number of factors including:

- the (g*A) value of glazed surfaces
- the absorptivity of opaque surfaces and the fraction of that energy transmitted into the building
- the multiple surfaces orientations
- the shading from nearby buildings and vegetation

B.1.1 Longwave radiative losses

In eq. B.6 the net radiative heat transfer Q_{sky} (W/m^2) between a body and the sky is given as a function of the net radiative fluxes and the view factor F_{sky}, which describes the proportion of radiation emitted by the surface which is received by the sky. ϵ_s is the long-wave emissivity of the surface, σ_B is the Stefan-Boltzmann constant (5.6697 x 10^8 W/m^{2}K^4) and T_s is the surface temperature (K).

$$Q_{sky} = \epsilon_s \sigma A_s F_{sky} (T_s^4 - T_{sky}^4) \tag{B.6}$$

Sky temperature T_{sky} (K) and emissivity ϵ_{sky} are not included the the weather datasets, but can be calculated as follows. From Stephan-Boltzmann law, the T_{sky} is related to the longwave global horizontal irradiance GHI_{IR} as eq. B.7, (Nave 2005).

$$T_{sky} = (GHI_{IR}/\epsilon_{sky}\sigma)^{1/4} \tag{B.7}$$

where Stephan-Boltzmann constant $\sigma = 5.6697 \times 10^{-8} W m^{-2} K^{-4}$. The sky emissivity can be calculated from available weather variables using the following approximation (Davies 2004):

$$\epsilon_{sky} = (0.787 + 0.764 \ln \frac{T_{dew-point}}{273})(1 + 0.0224N - 0.0035N^2 + 0.00028N^3) \tag{B.8}$$

where N is the opaque sky cover in tenths. Cloud cover C from weather files is in percent (%), so $N = C/10$.

B.2 Air exchange

Ventilation heat loss, through the exchange of warm internal air for cold external air, is thought to account for approximately 20% of a dwelling's heat loss (Cooper and Palmer 2013). Outside airflow into the building is mostly driven by pressure differences caused either by the winds or resulting from the difference between inside and outside temperatures. The heat loss due to infiltration is given in terms of volume flow rate q_v (m3/^{s}) and density ρ (kg/m^3) (1.205 kg/m^3 at standard temperature and pressure) (eq. B.9) and is due to the energy content of the exchanged air, defined by its the heat capacity c_p (J/kgK) (1005 J/kgK at standard temperature and pressure).

$$\Phi_{ve} = q_{ve} c_p \rho (T_{in} - T_{ex}) \tag{B.9}$$

B.2.1 Stack effect

The stack effect is the flow of air through a building resulting from a pressure differential created by the difference in temperature between internal and external air. The stack induced flow rate is a function of the internal/external temperature differential, building height, internal geometry, and openings such as vents, chimneys, unintended gaps and cracks.

An extensive theoretical treatment of these is given by Walker and Wilson (1998)

A convenient formulation of stack pressure in terms of internal and external temperatures (which are readily measured) is given by eq. B.10, where h is the building height, g is the gravitational constant and ρ_{air} is the air density of outside air (Walker and Wilson 1998) for the case where $T_{ex} \leq T_{in}$.

$$\Delta P_s = \rho_{air} g h (1 - \tfrac{T_{ex}}{T_{in}}) \tag{B.10}$$

Total infiltration is best approximated as using a power law eq. B.11 where C is the building leakage coefficient related to the area of leak openings, f_s is the stack factor, n is the flow exponent ($0.5 <$ n < 1) characteristic of the flow regime, and ΔP is the pressure gradient across the opening. For a typical residential building n ~ 0.67 (Walker and Wilson 1998).

$$q_{stack} = C f_s (\Delta P)^n \tag{B.11}$$

Substituting eq. B.10 into the volumetric flow rate eq. B.11 gives:

$$q_{stack} = C f_s \left(\rho_{air} g h (1 - \frac{T_{ex}}{T_{in}}) \right)^n \tag{B.12}$$

B.2.2 Wind driven ventilation

Wind induces a pressure on building surfaces eq. B.13, where S_w is the wind shelter coefficient - this has been expressed in various ways by different authors but the fundamental relation between wind induced pressure ΔP_w and wind speed v remains the same (Walker and Wilson 1998; Younes, Shdid, and Bitsuamlak 2012).

$$\Delta P_w = \tfrac{1}{2} \rho_{air} S_w^2 v^2 \tag{B.13}$$

This pressure induces a wind infiltration q_w, given by eq. B.14 where f_w captures the distribution of infiltration across walls, openings, flues, etc; and C_{wind} is another scaling factor (Walker and

Wilson 1998).

$$q_{wind} = Cf_w(\Delta P)^n \tag{B.14}$$

Substituting the wind pressure ΔP_w into the flow equation gives the wind induced flow rate eq. B.15.

$$q_{wind} = Cf_w(\tfrac{1}{2}\rho_{air}S_w^2 v_{wind}^2)^n \tag{B.15}$$

B.2.3 Total air change rate

Both stack and wind air flow are related to the pressure delta by a power law with common exponent n and flow factor C which relates to the total ventilation openings, and stack and wind specific factors f_s and f_w which can model the building's response to wind and stack effects. These can be combined in quadrature to give the total ventilation flow q_{ve} (Stamp, Lowe, and Altamirano-Medina 2013).

$$q_{ve} = \sqrt{q_{stack}^2 + q_{wind}^2} \tag{B.16}$$

Substituting eq. B.12 and eq. B.15:

$$q_{ve}^2 = (Cf_w(\tfrac{1}{2}\rho_{air}S_w^2 v_{wind}^2)^n)^2 + (Cf_s(\rho_{air}gh(1 - T_{ex}/T_{in}))^n)^2$$

This equation can be simplified considerably by lumping terms which are constants for a given dwelling, resulting in a wind coefficient C_{wind} and a stack coefficient C_{stack} (eq. B.17) - useful in the inverse modelling case where values for the separate constants may not be derived.

$$q_{ve}^2 = (C_{wind}v_{wind}^{2n})^2 + (C_{stack}(1 - T_{ex}/T_{in})^n)^2 \tag{B.17}$$

This model can be further approximated by a linear sum of wind and stack components (eq. B.18), with three constants to be determined for a building (Deru and Burns 2003).

$$q_{ve} = C_{ve} + C_{wind}v_{wind} + C_{stack}(T_{ex} - T_{in}) \tag{B.18}$$

Substituting into the heat flow equation eq. B.9 gives:

$$\Phi_{ve} = c_p \rho_{air} (T_{ex} - T_{in})(C_{ve} + C_{stack}(T_{ex} - T_{in}) + C_{wind} v_{wind}) \tag{B.19}$$