

Disaggregating High-resolution Gas Metering Data Using Pattern Recognition

Ammar Alzaatreh*, Lamine Mahdjoubi, Bill Gething, Francisco Sierra

Centre for Architecture and Built Environment Research (CABER), Faculty of Environment and Technology, University of the West of England, BS16 1QY, Bristol, UK

* corresponding author: ammar.alzaatreh@uwe.ac.uk

Abstract

Growing concern about the scale and extent of the gap between predicted and actual energy performance of new and retrofitted UK homes has led to a surge in the development of new tools and technologies trying to address the problem. A vital aspect of this work is to improve ease and accuracy of measuring in-use performance to better understand the extent of the gap and diagnose its causes. Existing approaches range from low cost but basic assessments allowing very limited diagnosis, to intensively instrumented experiments that provide detail but are expensive and highly disruptive, typically requiring the installation of specialist monitoring equipment and often vacating the house for several days. A key challenge in reducing the cost and difficulty of complex methods in occupied houses is to disaggregate space heating energy from that used for other uses without installing specialist monitoring equipment. This paper presents a low cost, non-invasive approach for doing so for a typical occupied UK home where space heating, hot water and cooking are provided by gas. The method, using dynamic pattern matching of total gas consumption measurements, typical of those provided by a smart meter, was tested by applying it to two occupied houses in the UK. The findings revealed that this method was successful in detecting heating patterns in the data and filtering out coinciding use.

Keywords: gas disaggregation, smart meters, gas heating, Building Physics, dynamic models

1 Introduction

1 The building sector holds great potential for saving energy and reducing CO₂ emissions, due to its large share
2 of global energy use and the relative cost-effectiveness of energy saving measures [1]. In order to tap into
3 this potential, policy makers across the world are imposing increasingly stringent requirements on building
4 energy performance and introducing schemes to incentivise energy saving measures. However, the actual
5 savings realised as a result of implementing those policies have repeatedly been shown to fall short of
6 theoretical predictions [2]. As the scale and extent of this performance gap [3] has become clear, policy
7 makers and the construction industry have become increasingly concerned that the shortfall could
8 fundamentally impact the ability of the building sector to deliver its anticipated share of the national carbon
9 reduction plan.

10
11 Addressing the performance gap requires methods to assess and understand the energy performance of
12 existing buildings at scale, using methods that allow diagnosing the reasons for the shortfall, whether
13 relating to physical or human factors. Such methods do exist, but as the existing building stock is largely in-
14 use, the challenge is for these methods to be applicable to occupied buildings with minimal cost and
15 disruption to occupants, but this issue has remained largely unsolved.

16
17 The main building performance measurement methods available are either based on intrusive heating
18 experiments, such as the co-heating test and its various similar alternatives (e.g. QUB [4], ISABELE [5]); or on
19 steady-state analysis of data recorded during normal operation and averaged over days or more, such as
20 Energy Signature methods and a host of linear regression-based ones [6, 7, 8, 9], which try to quantify the
21 relationship between heating energy and external-internal temperature difference and use it to characterise
22 the energy efficiency of the building.

23
24 The first approach can be robust and provide accurate and detailed outputs, but is highly intrusive and
25 expensive; it requires the building to be empty and unused for an extended period, along with extensive

26 instrumentation [10]. The second approach, in contrast, is more practical and much less intrusive, but typical
27 steady state models have long been deemed difficult to rely on for a detailed measurement that
28 distinguishes between the underlying causes as required [11, 12, 13, 14]. There are many reasons for this, of
29 which the most important are statistical limitations regarding correlation of inputs, as well as the reliance on
30 an assumption of consistent heating controls and occupant routines throughout the testing period. As such,
31 some methods of this class produce outputs characterising the energy performance of the building as it is
32 being used (e.g. the effective U-value, or the energy signature), which is no doubt an informative output, but
33 it is not intended to make a distinction between the effects of system efficiency, operation, and envelope
34 performance, and could well produce different ratings for physically identical buildings being used
35 differently [7], which makes it of limited use for diagnosing underlying issues affecting measured
36 performance.

37
38 The challenge is to have robust, detailed measurements of an occupied house, without disrupting
39 occupancy. Dynamic analysis of energy and environment data has long been suggested as a solution to this
40 problem [11, 15], and a growing body of research has demonstrated that it is potentially capable not only of
41 characterising the energy performance in terms of its causes, but also of attributing the fabric performance
42 to specific physical phenomena [16], in addition to having applications in optimal control [17], electric
43 demand-side management [18], and other forms of intelligent applications. However, this comes at the
44 expense of significantly increased instrumentation requirements compared to simpler options, of which the
45 most important is heating energy measurement [19]. Unlike steady-state methods, for which the total
46 energy consumption over long periods is usually sufficient, for dynamic analysis to be possible, it is
47 imperative that a high frequency time series record of the energy output of the heating system is obtained,
48 or, failing that, the energy used specifically for space heating. This is not as simple as making estimates by
49 subtracting the average consumption of other appliances, or using correlations with weather or water
50 consumption, as it has to provide the correct value at time steps traditionally less than 15 minutes. So far,
51 dynamic analysis studies in the literature have largely tended to target testing facilities or highly
52 instrumented buildings, like test boxes or office buildings with BMS systems installed [20, 19, 21]. When
53 these types of studies are attempted in more common conditions, or in residential buildings, extra dedicated
54 sensors have to be installed to monitor the energy used for heating [22, 23] among other variables.

55
56 The difficulty is most pronounced in the case of gas heating, particularly in residential buildings, where
57 measuring the gas used by the heating system alone, or the heat output of the heating system, requires
58 installing new meters, which are much more difficult and expensive to install than the other required
59 sensors (e.g. temperature), and could require being installed within the piping itself. Even when the existing
60 gas meter readings are sufficient, installing meter reading and logging equipment is problematic, due to the
61 wide variety of meter types in use and the various limitations attached to them. In a recent survey of
62 Building Performance Evaluation professionals in the UK [24], the consensus of respondents was that gas
63 heating is the most difficult element to monitor in a domestic setting. As gas-fuelled heating supplies 42.2%
64 of the total heating demand in Europe, and 80% of the total heating demand in the UK [25], this is a
65 significant obstacle to any practical large-scale application of dynamic performance assessment methods,
66 and the potential impact of addressing it is very significant.

67
68 In recent years, smart meter deployments have presented a great opportunity for these methods to become
69 practical, by making the collection of meter readings easy and unintrusive. However, smart meter readings
70 are still only a record of the total fuel usage of the building, which for natural gas typically includes domestic
71 hot water and cooking in the UK, providing no immediate insight into the individual end use of heating, and
72 certainly not the high frequency measurements of it required for a dynamic method. This challenge, often
73 known as the disaggregation problem [26], or single-channel source separation [27], has been attacked in
74 various ways in many contexts, but in cases similar to this one where a high frequency temporal sequence is
75 required, dynamic pattern recognition has often been used with success [28]. Examples of this include the
76 separation of readings from other types of smart meters measuring water use or electricity; however, no
77 such method has been attempted for gas systems yet, where existing statistical models for disaggregation
78 aim for separating sums [29], rather than high-resolution time series.

79

80 Though the aim is similar here to the case for electric or water smart meters, the volume of use for different
81 gas appliances is highly variable, and appliances usually control the flow with a mechanical valve subject to
82 various temperature or even manual controls, which produces very different load patterns than those found
83 in other types of metering. In addition, the effects of overlap in the case of combination boilers are more
84 complicated, as gas consumption is not simply the sum of the use by two end uses, and needs to be
85 estimated in other ways.

86

87 This paper presents a method for disaggregating the heating load out of the total recorded gas load by
88 analysing activity in high frequency metering data, selecting candidate activity windows based on step
89 changes in gas flow, and comparing them to a known instance of heating using Dynamic Time-Warping
90 (DTW) as a similarity metric to detect probable patterns. The required input is simply gas measurements
91 themselves, but weather and internal environment data, as well as information about the installed systems
92 were used to validate the results in this paper. This approach was tested in a real-life case study, involving
93 two occupied domestic buildings, with a similar building fabric, but contrasting gas energy use.

94

95 The paper is organised as follows; Section 2 reviews previous work, Section 3 describes the available data
96 and case study. Section 4 discusses the features appearing in the usage profile and characteristics of the gas
97 time series, Section 5 presents the proposed solution, and Section 6 presents results and evaluation,
98 followed by discussion and conclusion.

99

100

101

2 Previous Work

102

103 Dynamic models of the thermal behaviour and energy performance of buildings first became a topic of
104 research in the 1970s [31], and have since then advanced significantly [32, 33, 34], finding use in demand
105 side management [18], smart thermostats [35], in-situ building performance measurement (e.g. ISABELE [5],
106 QUB [4]), and more. There are many advantages to using dynamic models rather than steady state ones For
107 building performance characterisation in general, such as the ability to assign energy use to different
108 variables with confidence, but for an occupied building, the key advantage is that a dynamic model can
109 identify the effects of operation (or other changes) directly, whereas a steady-state one would be blind to
110 them, as the data is averaged over long time periods by necessity [11, 32].

111 Most dynamic models require, as a minimum, a high-frequency time series of internal temperature and heat
112 flow into the building (produced by a heating system), in order to relate the amount of thermal energy
113 expended to the degree of increase in temperature, and thus reason about thermal characteristics. To
114 acquire the input heat data in an occupied building, the easiest option is to simply acquire access to an
115 existing Building Management System (BMS) or instrumentation output from the heating system itself (e.g.
116 [36]), but these are rarely available in a residential building. As a result, instruments for data collection in
117 this sector have to be installed. This can take several forms varying in complexity and depending on the type
118 of the heating system involved.

119

120 For gas boilers, installing a heat meter is perhaps the most preferable option, as demonstrated in [34] and
121 [37]. This consists of a flow meter and a pair of temperature sensors on the supply and return pipes, which,
122 combined, provide a measure of the thermal energy supplied to the building. However, installing a heat
123 meter in an occupied domestic building is by no means a trivial task; it is often required to manipulate the
124 insulation around the pipes, or in the case of more conventional in-line meters, the piping itself; and access
125 to this pipework may not be easily available without significant disruption to the occupants, who could have
126 to arrange to be in while work is carried out, have to clean cupboards to allow access to the pipework, and
127 potentially suffer disturbance to the heating system. As a result, the commonly accepted alternative is to
128 record the measurements of the existing gas meter, typically in volume units, and then convert them into
129 energy. This is less direct and prone to error from several sources, such as the variation in fuel energy
130 content and boiler efficiency; which are not likely to be very problematic, and meter reading inaccuracy;
131 which is more likely to be so. Metering errors can result from many issues related to the metrological
132 performance of diaphragm gas meters, such as inaccuracy at low flow rates, and gas volume changing with
133 temperature [38]. Field tests of installed gas meters often find a number of issues with them [39] [40].

134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187

Nonetheless, as this method bypasses the requirement for installing new meters and reduces cost and difficulty considerably, it is a frequently taken option in many cases, particularly when heating is the only end use of natural gas in the building (e.g. [42], [43]). However, where gas is also used for domestic water heating or cooking, this method is problematic, as other uses have to be accepted as errors in the measurement of heating energy, or else resort has to be made to one of the earlier methods requiring added instrumentation. Thus, a method for disaggregating meter readings is likely to prove useful, at the very least by reducing the error margin back to that of the meter itself, rather than both the meter and presence of non-heating usage.

Disaggregating single point metered resource consumption by end use is a recurring problem in the literature. In the context of electrical consumption, this is known as Non-Intrusive (Appliance) Load Monitoring (NI(A)LM), and it has been an active field of research starting in the early 1980s. [44] Though early work in NILM made use of several measures specific to electric loads (e.g. using both voltage and current loads [45], reactive and active power [46], ..etc) or of very high frequency readings on the order of several kHz [46], more recent work has focused on using data obtained by measuring a single quantity (real power) at realistic intervals of 1 second or more, using methods mostly based on event detection and pattern recognition [28]. This type of data is presumed identical to the output of a smart meter should its frequency be high enough, but most deployed smart meters store and upload readings at a maximum frequency of one per 15 minutes. Some researchers have tackled the problem at this frequency or lower [47], but most choose to assume that readings will be recorded locally, as most smart meters are in fact capable of taking a reading every second or less, usually only for transmission to the local monitoring console or similar device. A review of the specifications of some of the most common models being deployed in the UK supports this assumption [48, 49].

Most methods employed in this class of problems rely on pattern recognition of appliance signatures, which are collected in dedicated experiments where the output of each appliance is recorded in isolation. With these signatures in hand, disaggregation is carried out by detecting switching events in the load graph, and then finding the sequence of appliance switching that is most likely to create the most matching graph [50]. This concept is widely applicable to problems other than electric loads, such as water metering [51, 52], but due to differences in the particular switching behaviour of the various systems, developed solutions are not usually portable to problems other than those for which they are intended.

Though methods developed to use smart meter-like readings have found use in different fields dealing with similar problems, the problem remains unaddressed for gas readings. Although statistical approaches to gas disaggregation have been developed [26]; they are intended only for disaggregating total volumes over some time period, rather than usage in the single time step, as required for dynamic models. Some hardware solutions to the problem have been introduced, in the form of extra sensors to be added to the metering infrastructure [53, 54], but of course solutions requiring production and installation of hardware are less desirable and do not solve the problem at a scalable level. For meter reading-based disaggregation, the only relevant published work to this problem so far is that of Bacher et al. [33], who, driven by the same aim of modelling the thermal behaviour of buildings without dedicated and intrusive sensor installation, manage to separate the total heat load of a district heated house in Denmark, recorded with a heat flow meter (MJ/h) at a resolution of 10 minutes, into space heating and water heating. Although this is not a gas system, the objective, being energy data of the heating system alone, is the same. In that work, the researchers differentiate heat load from water heating load by comparing periods of typical operation to a period without occupancy but with space heating operating as usual. They observe that the heat load in the house under study is slow-changing, of low relative magnitude, and is running continuously all day; whereas water heating takes the form of short, large spikes in the load graph, at magnitudes much higher than those of space heating. Based on that observation, they design a non-parametric estimator to filter out spikes in the total load graph, with the remaining, smooth usage assumed to be that of the heating system.

Though it may seem reasonable to expect similarity between the profile of heat load drawn from district heating and that of gas flowing into a local boiler, and though combi boilers do produce spikes when heating

188 water for short periods; the main feature, that water heating causes distinctive spikes unlike those of space
 189 heating, does not seem to hold true for gas boilers as a general rule, and gas cooking is not a spike at all. A
 190 part of the issue is due to the different modes of operation, as it is quite rare for a UK house to be heated all
 191 day [55], and a load spike would be expected once heating starts.

193 It is thus apparent that disaggregation of gas metering high resolution time series by end use has yet to be
 194 demonstrated, and this paper attempts to address this gap.

195 3 Data

197
 198 The data used in this study consists of time series of gas meter readings, taken at 1-min intervals in two 3-
 199 storey, 3-bedroom, terraced houses in the south west of England (pictured in Figure 1), for a period of two
 200 months. Both houses are social housing properties built in 1971 using Wimpey No-fines Concrete, but later
 201 improved by fitting double glazed windows and installing external insulation. Combination (combi) natural
 202 gas boilers provide central heating and instantaneous domestic water for both houses, but only one of them
 203 uses gas for cooking. Table 1 provides further information, and Figure 2 shows the floor plan for both.

204

	House 1	House 2
Type	End terrace	Mid terrace
Floor area	55 m ²	
External walls	No-fines concrete (280mm) + Paramount lining (40mm) Recently externally insulated with EPS (110mm)	
Heat Loss Coefficient (SAP Calculation) ⁱ	720 WK ⁻¹	
Floor	Thermoplastic tiles, 9mm screed, 100mm concrete slab	
Roof	Pitched roof with trussed rafters, 100mm glass mineral wool on ceiling	
Windows	Manual Double glazed PVC windows (closed most of the time)	
Ventilation		
- General	Windows + trickle vents (all rooms)	
- Bathroom	Manual intermittent extract fan	
- trickle vents	Above windows (closed)	
- kitchen	Extract hood	-
Hot water appliances		
- Shower	Electric Instantaneous	Gas
- Taps	Yes (gas)	
- Bath	-	Yes (gas)
Boiler model	Worcester Greenstar 25si	Vaillant EcoTECpro 24
- Type	Combi	Combi
- Max output	24kW	24kW
- Max DHW flow rate	10.2 dm ³ /min	9.6 dm ³ /min
Gas cooker	Yes	No
Heating control	Single thermostat in the hall, manually operated. TRVs in rooms. Mechanical timer control (unused).	Single thermostat in the hall, manually operated. TRVs in rooms. Time programmer (unused).
Secondary heating	Fan heater in main bedroom, 1500W	Portable electric heater, 2000W
Occupants	2	1

Occupancy	Employed (late hours, evening and night) + Student	Employed (early morning start)
------------------	---	---------------------------------------

205
206

Table 1: Information on studied houses



Figure 1: building with house 1

207
208
209



210
211
212

Figure 2: Layout of both houses

The owner (and main occupant) of house 1 is employed, working late hours usually starting late in the afternoon and returning at midnight, but they also work periodically overnight, leading to their hours of home activity mostly occurring around midnight or mornings. The other occupant is a student who is very often not present during the day, returning only in the evening. When the main occupant is not present, he tends to use heating for longer periods, and sometimes raise the thermostat set point. Heating seems to be manually switched on and off as needed, usually for one or two periods a day when the house is not empty, but the thermostat is rarely adjusted off its usual set point of 23°C, and the time programmer is never used.

The occupier of house 2 is in full time employment, with work hours of 6-2 most of the week. They seem to keep heating running constantly all year, even when outside the house or asleep, and even during summer time (apart from rare hot spells). The set temperature of the thermostat is often changed in winter months, varying between 25-30°C, but is left at 25°C in other times.

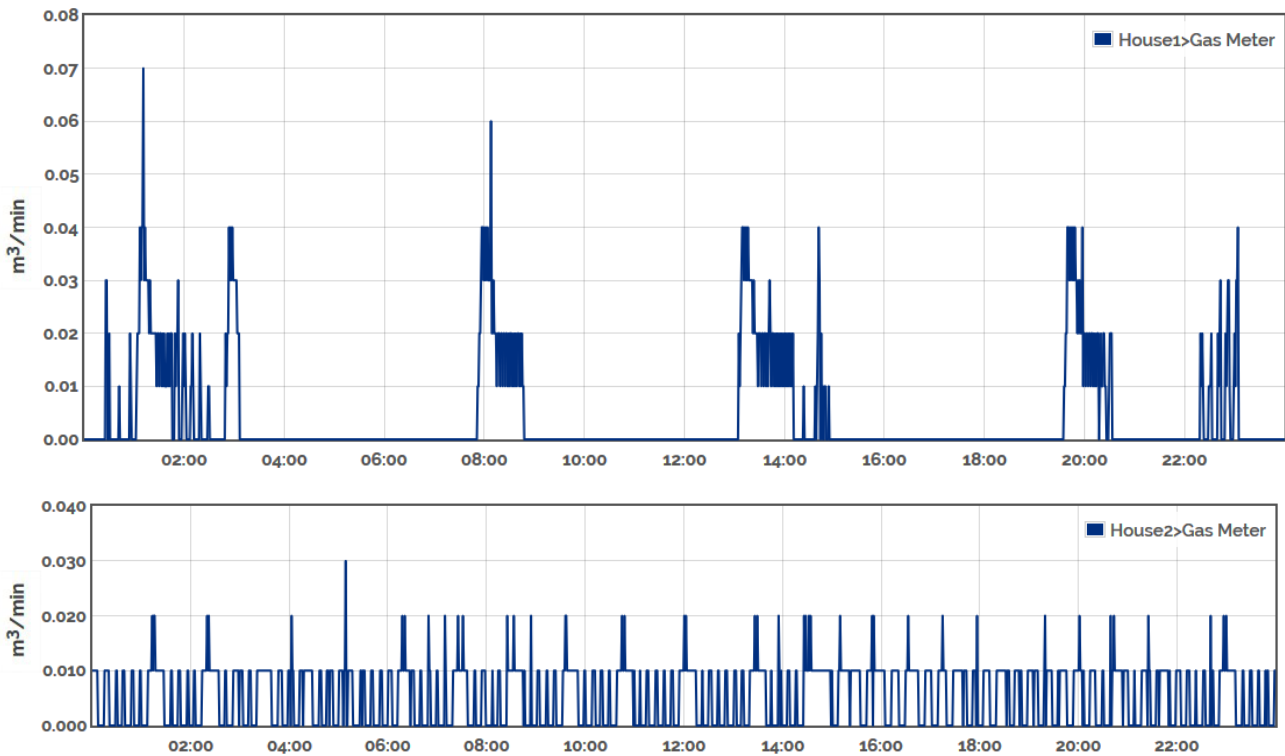
225
226
227
228
229
230
231
232
233
234

Gas usage is recorded using magnetic detectors installed on the available gas meters (Figure 3), which count the number of digit revolutions in the front of the meter (corresponding to 0.01m^3 of consumption, or approximately 0.11kWh of energy), and send the total count (in the last minute) to the data logger, every minute. The resulting readings can be interpreted as the increase in meter reading in the last minute (volume), or else as a reading of gas flow, in units of m^3/minute and resolution of 0.01m^3 , and can be converted into energy values by multiplying with the calorific value of the natural gas supply, which is given by the supplier as $39.4\text{ MJ}/\text{m}^3$. Figure 4 shows an example of the output for a 24-hour period in January 2017.



Figure 3: sensor for meter reading

235
236
237



238

Figure 4: 24 hours of gas usage as recorded by the sensors in both houses

239
240
241
242
243
244
245
246
247
248
249
250
251
252
253

The monitoring frequency of gas usage, at 1 reading per minute, is chosen to capture the dynamic variation in gas use in as high a granularity as can be done while still being within the expected practical capabilities of a smart meter-based solution for large-scale deployment. That said, these readings are not truly representative of a smart meter reading in terms of accuracy, as the metrological performance of gas smart meters is expected to be superior to older types, simply by virtue of being newer, or potentially due to using newer technologies, such as ultrasonic (e.g. [40]) and thermal mass measurement, in some jurisdictions.

In the present study, the specifications for the available meters are given in Table 2, together with the calculated average flow rate in both houses. Both meters are equipped with temperature compensation elements, and have a magnetic pulse rate of 0.01m^3 , meaning that the dials only trigger the magnetic detector at steps of 0.01, despite the resolution of the meter reading being well below that.

	House 1	House 2
--	---------	---------

Model	Itron U6	Elster BK-G4
Max flow rate	6 m ³ /h	6 m ³ /h
Min flow rate	0.12 m ³ /h	0.04 m ³ /h
Cyclic volume	2 dm ³	2 dm ³
Magnetic pulse rate	0.01 m ³	0.01 m ³
Reading resolution	0.0002 m ³	0.0002 m ³
Average flow rate in house (when in use)	0.18 m ³ /h	0.3 m ³ /h

Table 2: Gas meter specifications

254

255

256

257

258

259

260

261

262

263

264

265

266

267

268

269

270

271

272

273

274

275

276

277

278

279

280

281

282

283

284

285

286

287

288

289

290

291

292

293

In addition to gas meter readings, wireless air temperature sensors were placed in every room, and a weather monitoring station was installed on top of a nearby building to record local weather variables. This data was used for the purpose of development and validation, but is not part of the disaggregation method itself, and not required for application.

4 Signature

The basic premise of this work is that a heating system operates in a regular fashion, producing a temporal pattern of gas consumption sufficiently different from those of other systems that it can be relied upon to distinguish heating from other uses. Recognising and recording an appliance pattern could be done with a simple test, but unlike the electrical counterpart, there are usually only a few appliances using gas in a typical domestic building, and quite frequently, only one of them would be in use. This eliminates the need for dedicated experiments or sub-metering to record a signature, as the operation of different appliances can generally be discerned quickly by the analyst; however, some of these signatures can change in an analog way in response to human control or temperature variables, as with gas hobs for instance, which means defining a signature is not always simply a matter of extracting a 'clean' instance of it, nor of defining several states or stages.

For the present purpose it suffices to find the approximate signature(s) of the heating system, as all other usage will be discarded, and this signature is usually less difficult to work with. The operation of the modern gas boiler is largely dependent on a modulated firing cycle, switching between several levels of intake valve opening in response to the current temperature level of the fluid or a defined operation mode. This can produce several possible signatures, but the general shape is consistent between most boilers, composed of a high initial load as the boiler brings the heat up from rest, followed by modulated operation, which turns gas firing on and off, or high and low, as needed to preserve a certain target level of temperature. When the fluid temperature is already relatively high, the initial high load may not occur, and modulation may produce lower, or shorter pulses, depending on the boiler make and configuration. The result can take the shape of a train of slow on/off cycles, or a fast oscillation, or several stages of multi-level firing, but is in all cases a predictable effect of modulation.

Identifying the signature visually in the total load graph is not a difficult task. In most cases, heating operation will have a consistent daily or weekly pattern of occurrence, is absent in warmer days, and makes up the largest volume of gas use of all gas appliances, which makes it easily detectable on inspection. Any of these properties could possibly change depending on the type of system and the way it is operated, but save for some edge cases, this pattern is always likely to be the most prominent one. Figure 5 shows a 24 hour sample of the data in house 1 as an example, together with internal temperature data for validation.

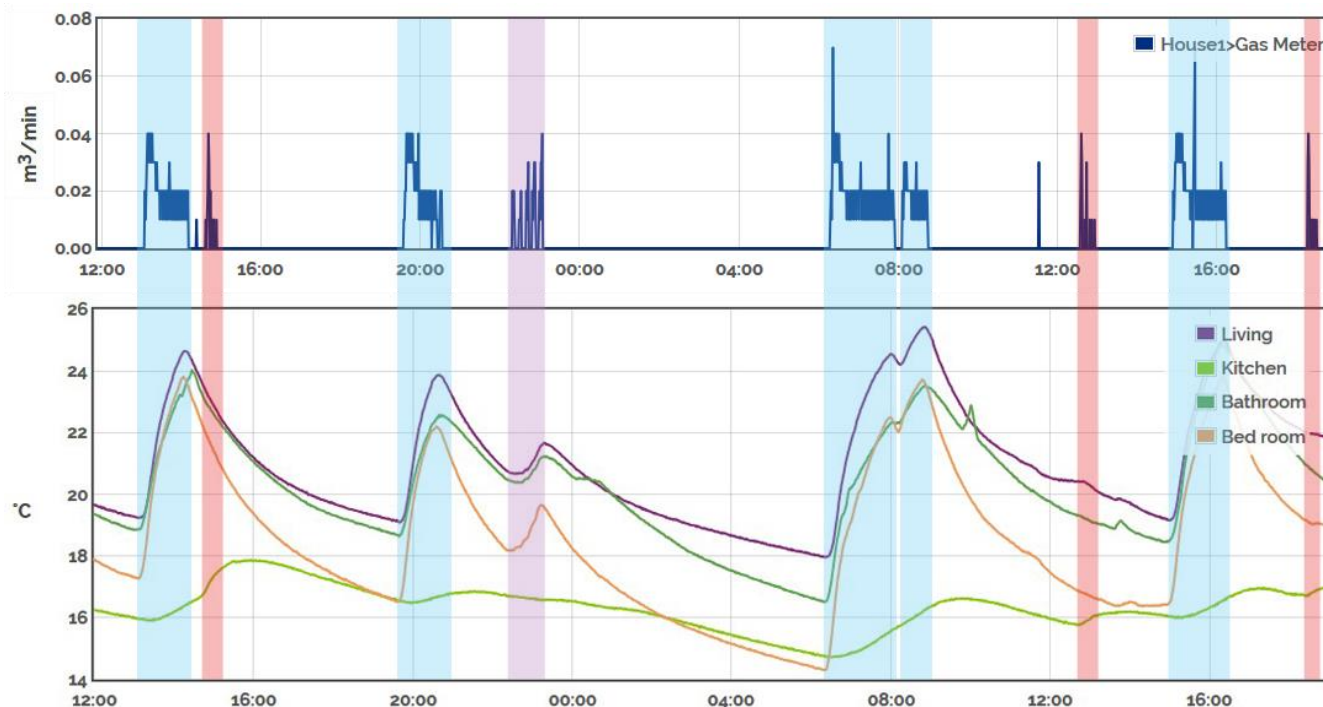


Figure 5: minutely gas usage graph for House 1 with patterns highlighted

294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312

Looking at Figure 5, it can be seen that three distinct patterns are immediately observable in the graph for House 1, with minor variations to each and some spikes. The most obvious candidate for heating is the first occurring pattern, being the most prominent and largest one, and matching the modulation effect expected of heating. This can easily be verified by observing temperature changes in the house at the same times. The pattern repeats five times in this period, and in all cases causes a sharp rise in the temperature in all rooms, which confirms that it is indeed caused by space heating. At approximately 13:30, a different pattern can be seen causing a rise in the temperature of the kitchen, but none anywhere else, which indicates that it belongs to gas cooking. The third pattern, at approximately 23:00, also causes a rise in temperature everywhere, but in a much slower manner than the first one. This is certainly also a result of heating, likely when less heating power is required. The remaining part is the irregular spikes, which are a typical result of water heating in a combi boiler.

Thus we can see that for House 1, the signature takes the form of a sudden rise in gas flow, peaking at $0.04\text{m}^3/\text{minute}$ (26.2 kW) and then falling again, over approximately 20 minutes. This is followed by a steady flow with a small oscillation, which continues at the same rate until heating is turned off.


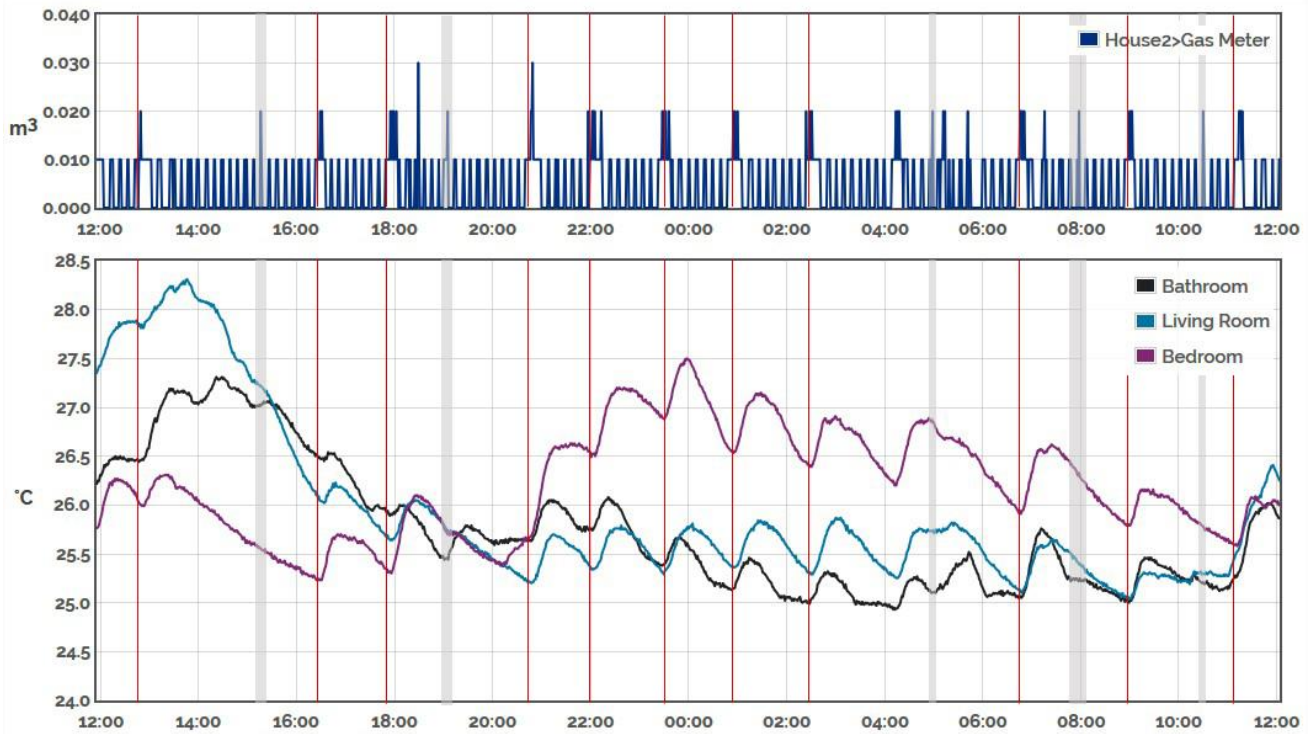


Figure 6: Space heating gas usage patterns in house 1

313
314
315
316
317
318

In House 2 heating is left running continuously, as this house is kept at a constant, high temperature all day long (rather unusual for a UK dwelling), but several occurrences when heating was turned off for several


319 hours are available in the data. The observable signature during normal operation, as seen in Figure 7, is
320 simply a series of rectangular pulses at a constant flow level but modulated width. Periodically, as
321 temperature falls, this is replaced with a longer period of constant flow at the same level, with a spike or
322 several spikes to a higher one (crossed with vertical lines in the figure).



324 **Figure 7: Gas usage patterns and temperature in House 2**

325
326
327 As temperature rises clearly coincide with the periodic spikes, we can be confident that they are caused by
328 heating. Spikes that do not occur during the longer flow period, like those observed at 15:10 and 18:20 in
329 Figure 7 (highlighted), do not seem to coincide with a temperature rise, or else with a small rise in bathroom
330 temperature, which likely indicates water heating for a tap load.
331 The remaining significant patterns that can be observed in the data are shown in figures 8 and 9. Following
332 the same reasoning process, these can be attributed with confidence to heating starting from rest and
333 longer water heating loads, respectively.

334



335
336
337

Figure 8: Heating gas usage patterns in House 2




Figure 9: Gas usage pattern for extended water heating (e.g. running a bath) in House 2

338
339
340

341 It is evident that space heating signatures in both houses are discernible simply by observing the time series
342 of measurements and identifying the most prominent pattern with a large volume and regular recurrence
343 even in the absence of temperature data and information about the ratings or particular make and model of
344 the boiler. In the process of conducting this study, the information used for validation was only consulted
345 after signatures were initially identified. Expectation of the possible shapes for the profile of usage was
346 guided by an understanding of the operation of the installed gas boilers, but the main reliable characteristic
347 to look for was simply the large volume and regular recurrence.

348
349

350 **5 Method**

351

352 Based on the observations in the previous section, it seems to hold true that heating has a relatively
353 consistent signature, which makes pattern recognition an obvious candidate for the task of detecting it
354 among other uses.

355 The objective of the presented method is to detect periods when heating is on and filter out any non-
356 heating usage, arriving at a record of the gas used for heating per minute. This is achieved by comparing the
357 gas usage profile in all time windows of non-transient activity to all defined space heating signatures using a
358 time series similarity metric, selecting windows that have a similarity above a certain threshold, and filtering
359 them to remove any added usage that deviates too far from the average of all windows. A summary is
360 shown in the flow chart in Figure 10.

361 In presenting this method we assume that the signature of the heating system has already been identified,
362 whether by visual inspection, experimentation, or one of the other ways discussed in the previous section.

363

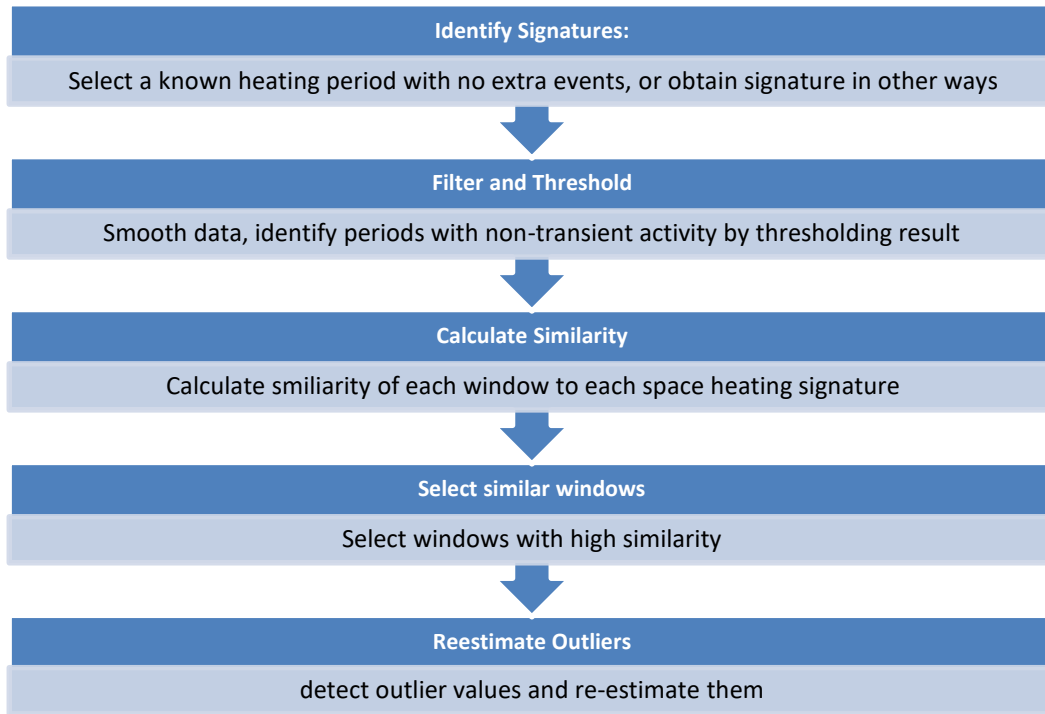


Figure 10: Flowchart of method

364
365
366
367

368 5.1 Window detection

369

370

371

372

373

374

375

376

377

378

379

380

The first step is to select candidate windows of activity, where a candidate window is one with continuous, sustained usage, beginning with a significantly large positive event (increased flow) and ending with a significantly large negative event (decreased flow). To select these windows, we begin by applying a smoothing filter to the time series, which creates a copy in which spikes and oscillations are removed, and all changes in the graph are smoothed out, simplifying it into coherent windows of operation (figure 11). Selection of candidate activity windows can then be carried out using the filtered version, simply by imposing a threshold value, and selecting any window crossing this threshold.

Starting with a series of gas readings $\{g_t\}$ taken at regular intervals, and the symmetrical moving average as a typical choice of smoothing filter, the series of smoothed readings $\{G_t\}$ can be calculated by

$$G_t = \frac{\sum_{t=-(S-1)/2}^{t=(S-1)/2} g_t}{S} \quad (1)$$

381

382

383

384

Where S is the length of the averaging window.

Using the filtered output, window boundaries can be found by matching threshold crossing events in the smoothed series, for a certain threshold value of T .

$$\delta_t = \begin{cases} 0, & G_t < T \\ 1, & G_t \geq T \end{cases} \quad (2)$$

385

$$\Delta\delta_t = \delta_t - \delta_{t-1} \quad (3)$$

386

387

388

389

390

391

392

Where $\{\Delta\delta_t\}$ is the time series of threshold-crossing (switching) events, indicating a rise in flow for positive events ($\Delta\delta_t = 1$), a fall for negative ones ($\Delta\delta_t = -1$), or no change ($\Delta\delta_t = 0$). The smoothing of the series, if calibrated correctly, guarantees that these events will not be spikes or fast perturbations, but the start or end of sustained activity of a significant time length in most cases.

To define candidate activity windows, event times are split into positive and negative series. So for every $t_n \in t$, instances are distributed according to

393

394

$$t_n \in \begin{cases} E^+, & \Delta\delta_{t_n} = 1 \\ E^-, & \Delta\delta_{t_n} = -1 \end{cases} \quad (4)$$

395

396

397

398

399

Assuming g starts and ends at 0 (or is padded with zeros), both E^+ and E^- will have the same length, meaning every rise above the threshold will be followed, eventually, by a fall below it. Thus candidate activity windows are simply subseries of the original gas readings series, between each two subsequent threshold-crossing events.

$$W_n = \{g_t\}_{t=E_n^+}^{t=E_n^-} \quad (5)$$

400

401


402

403

404

405

Arriving at this result by following these steps requires a choice of two parameters; the span of the moving average window S and the threshold T , which have appropriate values different for each setup, and possibly for each signature. In this case they were tuned manually, but should the average length of heating activation in the house be known, they may be optimised to yield a match to the average candidate window length or some similar condition, but this is left for future work.



406

407

408

409

410

411

5.2 Similarity

412

413

414

415

416

417

418

419

420

421

422

423

424

425

426

427

428

429

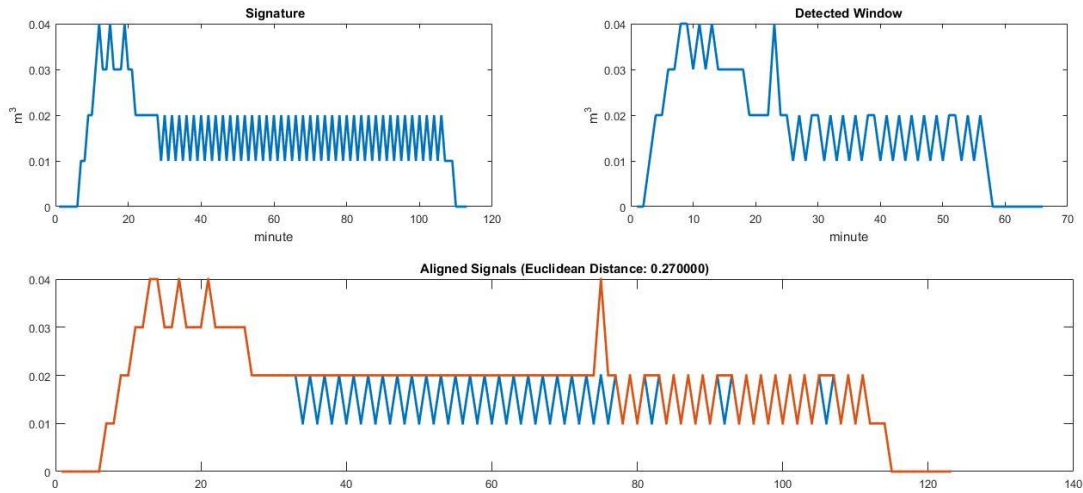
430

431

To distinguish heating from other uses, it suffices to detect the signature patterns. This is done by comparing identified windows of activity with the signature, using a standard similarity calculation algorithm. In the field of time series data analysis, particularly for the purposes of classification and clustering, Euclidean distance is widely used to measure the similarity of series, but its sensitivity to time axis distortions is a well-known problem [56], wherein slight changes in the timing of a certain feature (e.g. peak) or the length of time between features within the series would affect the measured similarity highly. This has led to the introduction of Dynamic Time Warping distance (DTW) [57], an alternative similarity metric that has proven very highly successful in many fields, such as speech recognition [58] and bioinformatics [59, 60], and has been successfully applied to disaggregation problems similar to the present one, in both electricity [61] and water [51].

Dynamic Time Warping measures the ‘distance’ between two time series irrespective of any variations in the time dimension. The algorithm attempts to temporally stretch or compress parts of the compared signal until it finds the state with the shortest Euclidean distance to the standard one and the same length, which in this case allows the comparison of windows to signature irrespective of variations in length of firing stages or modulation times, and using only the data points themselves rather than any model of them. A side effect of this is that it sometimes inflates the similarity scores of windows containing very short or very simple profiles if the warping path is not optimised correctly [62]. This is not a serious issue however, as a properly selected threshold for window selection should cause most of these to be dropped in the detection stage.

432 Figure 12 shows an example of DTW similarity computation. The best match between the detected
 433 candidate window and the signature is produced by stretching the detected signal to have the same length
 434 as the signature. It can be seen that the detected window seems to be an almost perfect match aside from
 435 an extra spike right before the modulation stage, which could be due to an overlapping use of gas, or due to
 436 heating itself, as will be discussed in the next subsection.



437
 438
 439 **Figure 12: Alignment between signature and a detected window**
 440

441 The DTW distance of all selected windows to the signature is a measure of similarity, on which we can
 442 impose a threshold to select windows that are sufficiently similar to the signature to be judged as heating
 443 windows, with a simple condition

$$444 \quad H = \{W_n \mid DTW(W_n, Sig) \leq \theta\} \quad (6)$$

445
 446 The set of heating activity windows includes every window W_n in the set of all candidate windows W where
 447 the DTW distance from a certain signature is less than a threshold θ .
 448
 449

450 5.3 Anomaly filtering

451
 452 Although heating is usually the dominant end use in terms of total time and volume, other uses may still
 453 coincide with a heating event in the same time window. In the absence of such events, the observable
 454 profile of heating in the time series matches the stored signatures quite closely, but variations and
 455 additional spikes do occur. The problem of filtering, thus, is to identify which differences from the signature
 456 are anomalies, resulting from non-heating usage rather than a natural variation in heating operation, and
 457 where this is determined to be the case, estimate the heating-only reading.
 458

459 Anomaly detection is a common problem in data mining, with many existing solutions. Simply stated, the
 460 aim is to identify events that do not conform to expectation, whether the expectation is explicit (based on a
 461 known model, stochastic distribution, or known sequence of normal behaviour) or has to be inferred from
 462 the data itself [63]. The key insight relating to this case is that non-heating end uses are unlikely to highly
 463 correlate with heating events in time, and do not occur consistently at the same stage of heating operation
 464 with the same magnitude, which means the unwanted variations are simply those that occur infrequently.
 465

466 In the context of the disaggregation problem, the signature is this 'expectation', often defined in terms of
 467 Markov chain models [28, 64], but in this case the signature is only a sample, rather than an accurate
 468 description or model of operation. This can be changed at this stage by training a model of detected heating
 469 windows, to serve as an expectation for anomaly detection. As the behaviour of heating is dynamic, relating
 470 to a stage of operation and values at several time steps in the past, and the set of possible values is limited,

471 we can define this using a matrix of multi-step (reverse) transition probabilities, as an extension of the
 472 markov chain model. This approach is common in anomaly detection of discrete sequences [65].
 473 Let X denote the state at time t , where $X_t \in \{1, 2, 3, \dots\}$ corresponding to the set of possible values of g_t ,
 474 we can have matrix P where

$$475 \Pr(X_t = i | X_{t-1} = j, X_{t-2} = k, X_{t-3} = l) = P_{i,j,k,l} \quad (7)$$

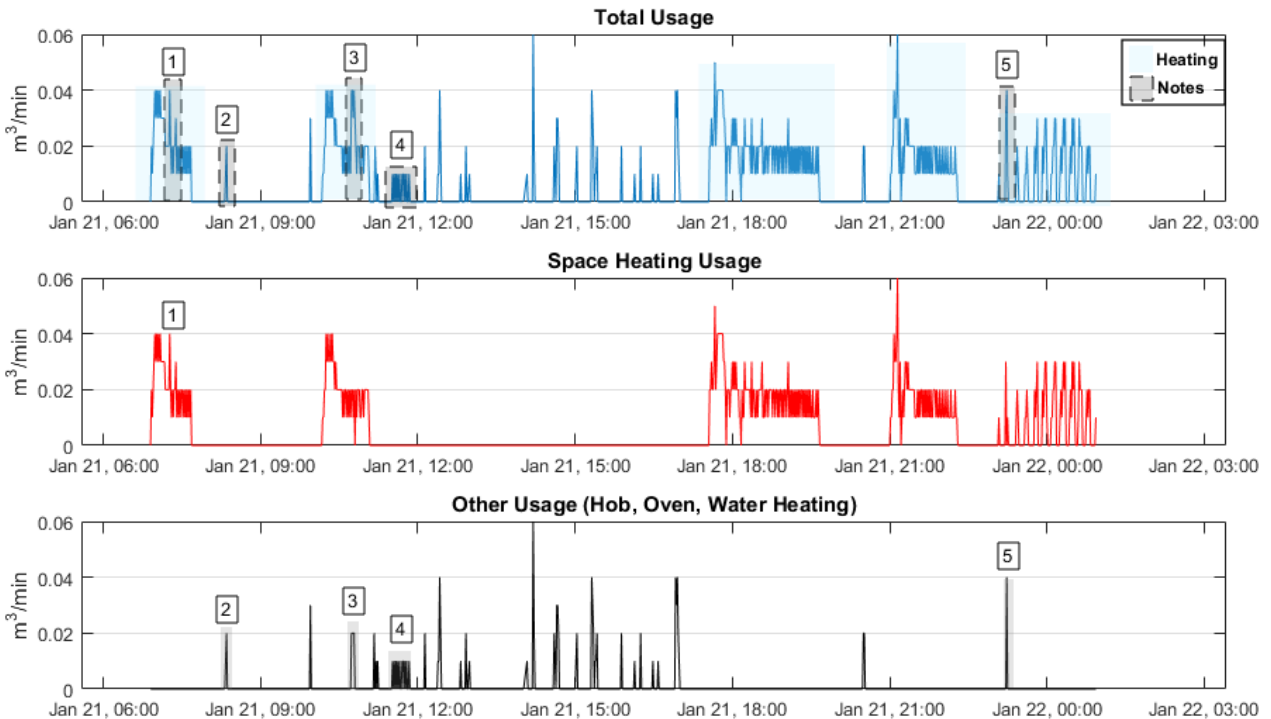
476 This matrix is easily populated by calculating the probability as the frequency of each value given the
 477 previous values, within the detected heating windows themselves, with points outside window boundaries
 478 assumed zero. This matrix must be calculated independently for each set of windows matching a certain
 479 signature, and then anomalies are simply any value where the probability is less than a certain threshold.
 480 When an anomaly is detected, a good estimate of the correct state at that point is simply the most probable
 481 state

$$482 \hat{X}_t = \operatorname{argmax}_x \Pr(X_t = x | X_{t-1}, X_{t-2}, X_{t-3}) \quad (8)$$

483 But for a combi boiler where only one use is activated at a time (space or water heating), this can simply be
 484 replaced with a value of zero.
 485 This method of anomaly detection and re-estimation is rather generic and possibly computationally
 486 intensive. It depends on the possible readings being a discrete set, which, although not an obstacle to their
 487 application where better resolutions are available, wastes the information possibly contained in them. Other
 488 more elegant methods can be used here, particularly robust estimators, like that shown in [33], but for the
 489 present purpose this is deemed sufficiently effective.
 490

491 6 Results

492 Applying this method to the available data yields a separation between space heating and other uses, as
 493 shown in figures 13 and 14. Figure 13 shows a day of results in House 1, the total usage graph (top), is split
 494 into a heating graph and a non-heating one by applying the methods described in the previous section. By
 495 visual inspection alone, it can be seen that all recognisable heating events have indeed been isolated, and
 496 several abnormal spikes in them have been smoothed out. Some features of interest are highlighted and
 497 numbered, with discussion to follow.
 498



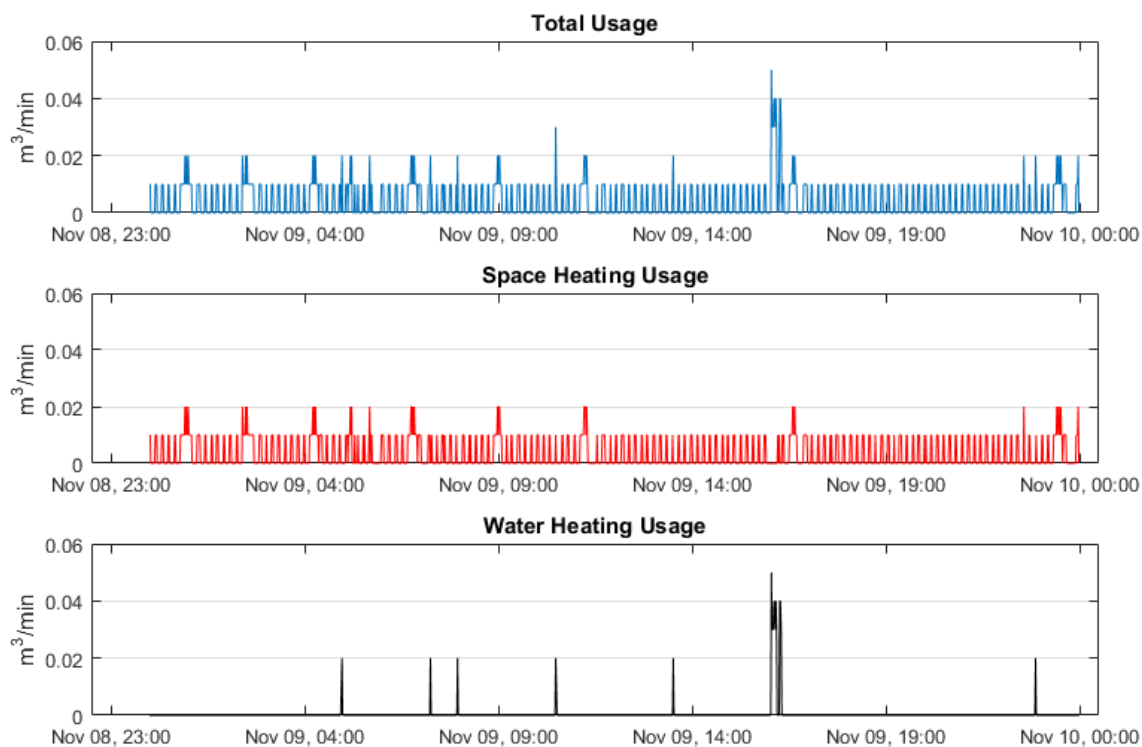
501
 502 **Figure 13: Demonstration of results in house 1**

503

504 Looking at the total usage graph, four instances of the main identified heating pattern can be discerned, and
505 one of the second one. Starting at 07:00, the pattern is a match for heating, but it differs from the signature
506 shown in Figure 6 by a spike in the middle (1). This spike has not been filtered out, as it happens often at this
507 point in the heating cycle, and is thus calculated to be a variation in boiler operation. Near 08:30, marked by
508 (2) is an isolated spike of gas usage, too short to be the result of space heating, and the same can be said of
509 the following one, around 10am. Both are likely results of water heating, and are both filtered out. Right
510 after, a heating cycle is detected, but in this case, after 30 minutes of uninterrupted heating, usage picks up
511 unexpectedly (3); as this is irregular, the calculated probability of it being a variation in boiler operation is
512 low and it is thus filtered out.

513 Shortly after 11:30 another pattern of usage appears (4). This pattern does not fit any heating signature and
514 is thus removed into the other usage bin. Given its low volume, consistent level, and oscillation profile, this
515 is likely to be gas oven operation. Thereafter, until 17:30, multiple spikes appear of different characteristics,
516 but none of them is long enough to be detected as heating, and are thus removed. At 21:30 (5), a low-load
517 heating pattern is detected, but a spike in usage a few minutes after it starts is judged irregular and
518 removed. It is notable here that the anomaly filter does expect a spike of $0.03\text{m}^3/\text{min}$, as evident by the
519 output, but the following increase up to $0.04\text{m}^3/\text{min}$ for 2 minutes was considered anomalous, showing that
520 even if non-space heating results coincide with space heating, they can be detected.

521 In House 2 the results are easier to explain. Figure 14 shows a full day of results. It can be seen that space
522 heating is running all day, causing a constantly oscillating gas flow, switching regularly back to a higher-use
523 pattern, which appears three times between 08:00 and 09:00. Spikes that do not fit this profile have been
524 filtered out. In particular, a period of extended use around 16:00 is removed as it does not fit the heating
525 signature. This pattern is likely to be a result of extended water heating use, possibly caused by running a
526 bath, as explained in section 3.
527



528
529 **Figure 14: Demonstration of results in house 2**
530
531
532
533
534
535
536

537 **7 Validation**

538

539 Correct validation of these results requires a ground truth series, where gas usage for each of space heating,
540 water heating, and gas cooking is measured independently. As the boilers in this case are both combi boilers,
541 both water and space heating make use of a single gas intake pipe, and independent measurement is not
542 possible to do externally. This, combined with occupancy restraints, makes correct direct validation difficult,
543 and recourse has to be made to less ideal options. In order to validate the results, three tests were carried
544 out. First, an experiment is carried out in House 2, in which water heating is triggered at set times to
545 interrupt space heating, and the method is used to detect this. Second, a model predicting temperature rise
546 in the houses based on gas usage is created, and then applied to show that disaggregated space heating gas
547 usage, as given by the method, predicts temperature rises more accurately than the original total. Finally,
548 the average daily profile of results is compared to the expected based on knowledge of occupant routines,
549 showing that the results are consistent with expectations.


550

551 **7.1 Experimental validation**

552 The possibility of full experimental validation was constrained by occupancy considerations and boiler types,
553 as explained above, but short experiments were possible in house 2, where gas is used only for space and
554 water heating.

555 In the first experiment, heating was turned off for 30 minutes, then started again from rest, and left running
556 for 65 minutes with no water heating usage. The result, shown in figure [15], is consistent with the expected
557 pattern found in section 4.

558




559

560 **Figure 15: Recorded gas usage during experiment 1**

561

562 In the second test, while space heating was left running, water heating was triggered for periods of 3
563 minutes and 2 minutes, with a pause of 10 minutes in between. This resulted in gas usage rising consistently
564 during these periods. By applying the method to the total usage data, the water heating usage is indeed
565 detected and removed in both instances, as shown in Figure 17.

566



567

568 **Figure 16 Total gas usage during experiment 2**




Figure 17 Disaggregated space heating usage during experiment 2

569
570

571

572 7.2 Temperature modelling

573

574

575

576

577

578

579

580

581

582

583

584

585

586

587

588

589

590

591

592

593

594

595

596

597

598

599

600

601

602

603

604

605

606

Detailed temperature data is available in the studied houses, and the effect of space heating on whole-house temperature is readily observable and distinguishable from that of other gas uses. This provides a highly suitable opportunity for validation of our method.

The most direct observable effect of space heating is, of course, a significant increase in whole-house internal temperature, whereas gas cooking and hot water heat is localised, and in the case of the latter, a significant portion of it is rapidly lost into the drainage system. Given a model of whole-house temperature change based on gas heating, these differences can be used to validate the results of the method, by comparing the accuracy of temperature predictions using total gas usage data, to those using disaggregated heating usage data. Correct disaggregation should lead to more accurate predictions, as non-heating gas usage would lead to erroneous predictions of temperature rise.

Dynamic modelling methods used to predict thermal response in buildings are numerous and varied, but for the present purpose, an autoregressive (AR) model is well-suited, as it encodes dynamic, short term response, without requiring a tailored model based on the physics of the building. AR models and their derivatives, such as the ARX (eXogenous) model, have been used to predict temperature and energy consumption in buildings in many studies under wildly different conditions in all parts of the world [36] [66], [67] [68].

The process for fitting an appropriate ARX model to energy and temperature data is well known and covered by several sources [69] [66], and is thus dropped here for brevity. The following model is found to be a good fit

$$T_i^t = T_i^{t-1} + b_1 \varphi^{t-1} + b_2 \varphi^{t-2} + \dots + b_5 \varphi^{t-5} + c_1 T_e^{t-1} + \varepsilon \quad (9)$$

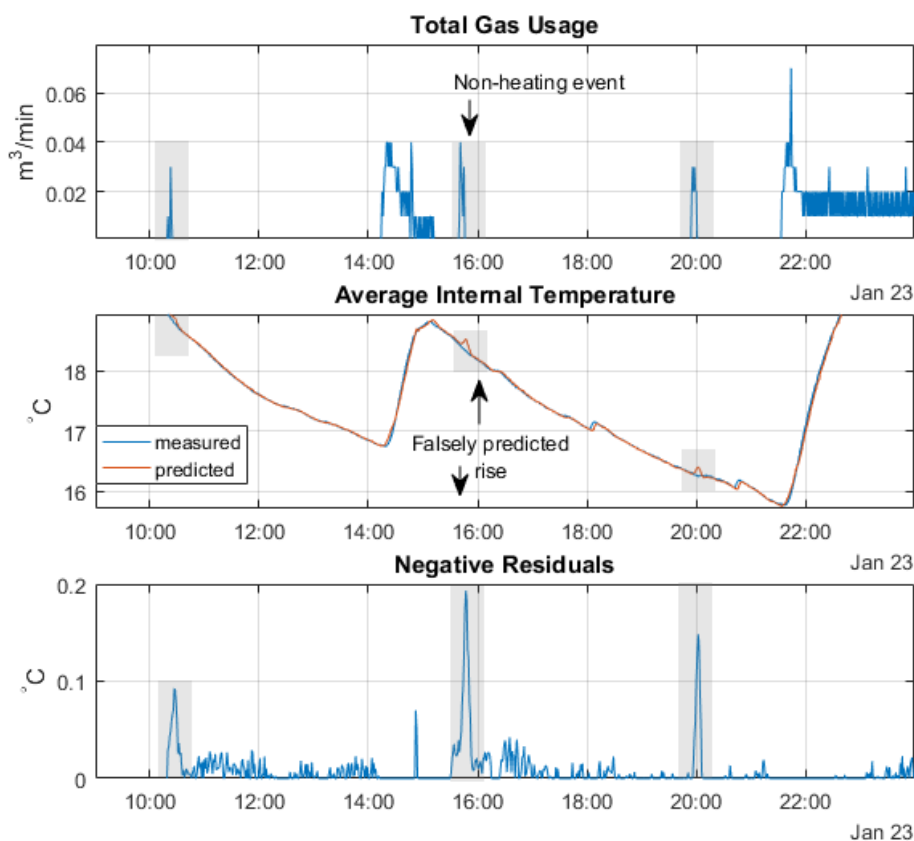
Where T_i^τ is the average internal temperature at time τ , T_e^τ is the external one, and φ^τ is the gas usage during time step τ . Effectively, this model predicts that the temperature at time t will decrease from that at time $t-1$ by an amount linearly proportional to the difference between the internal and external temperatures, while also increasing by an amount linearly proportional to the gas used for heating in the last 5 minutes. The factors are calculated by a statistical fit to the total available data, and the results of model optimisation, in 3 steps, are shown in table 3 below.

Metric	Value
1-step ahead prediction fit	97.82%
5-step ahead prediction fit	93.94%
Mean Square Error (MSE)	0.0011
Final Prediction Error (FPE)	0.0011

Table 3: Parameters of fit for ARX model

607
608
609
610
611
612
613

As the space heating load is the major part of gas usage in the data, the model is tuned to model the response to space heating, with some bias resulting from the presence of non-heating usage. When supplied with the original total gas data, before disaggregation, the results of 1 step-ahead prediction fit the actual measurements closely (94.12%); however, when a non-heating event occurs, the model erroneously predicts an increase in temperature. Figure 18 shows an example of this.



614
615

Figure 18 Errors in prediction caused by non-heating events. The last panel only shows negative errors

616 Using the output of the disaggregation method, the same model was used to predict temperature again, this
617 time with possible non-heating gas usage removed from the input. The results show a significantly better fit
618 to the measured temperature, and a significant decrease in the mean square error, (table 4). A comparison
619 of fit quality before and after disaggregation is given in table 4 below.
620

Metric	Value (total gas)	Value (disaggregated heating)
1-step ahead prediction fit	94.12%	97.82%
5-step ahead prediction fit	85.12%	93.94%
Mean Square Error (MSE)	0.0018	0.0011
Final Prediction Error (FPE)	0.0078	0.0011

621

Table 4: Parameters of fit for ARX models, using total and disaggregated heating data

622 Given that heating is the largest use of gas by a significant margin, these differences are quite significant.
623 This is a clear indication of the validity of the results. The quality of fit of the same model increases

624 remarkably after disaggregation, proving that at the very least, the disaggregated series contains
625 significantly less non-heating usage than the original one. Unfortunately, this method cannot be used to
626 detect whether all non-heating usage has been removed, as errors in prediction can stem from other
627 sources, such as unmonitored heat sources (e.g. electric heaters), changes in ventilation (windows,
628 mechanical ventilation, ..etc), and solar gains. For the same reason, using such a model for disaggregation in
629 the first place, though possibly effective, may introduce bias in the results of any analysis based on it, but
630 that question is still open for study.

631

632

633 This confirms the validity of the outputs, particularly as those instances were not easily detectable by sight,
634 but as the experimental data by no means covers all expected circumstances, this is not by itself complete
635 validation, and thus other validation tests are presented below. A more complete validation exercise would
636 be desirable for future work, but is not possible in the context of this study.

637

638

639 **8 Discussion**

640

641 Prevailing approaches for investigating the performance gap have had limited success so far in providing the
642 required diagnostic capabilities. At present, there are several challenges to the use of dynamic methods in
643 occupied buildings at the required scale, but data collection remains a key obstacle for research and actual
644 implementation, which has led to increased interest in simpler methods that lack accuracy and do not
645 provide the diagnostic insight required for tackling the gap.

646 The presented method improves the potential applicability of dynamic analysis methods to existing occupied
647 residential buildings by making dependence on smart meter data an option for heating energy measurement
648 even when other uses are present. This reduces the potential cost greatly, as any alternative requiring the
649 installation of extra instrumentation would entail the cost of the hardware and the labour and time of skilled
650 technicians; whereas this method only requires the addition of some processing step to the smart metering
651 infrastructure. This can take the shape of a small program running on the same device being used as an in-
652 home monitor, which cuts down the cost of deployment to those of a software update for the existing
653 monitors. There are already calls for smart metering standards and commercial devices to incorporate the
654 ability to host different software and forms of feedback, rather than simply rely on the software provided by
655 the utility company [66], this could pave the way to a different service model for employing smart meter
656 data, in which service providers can create applications that the consumers can give access to their data to,
657 similar to the Green Button Initiative model introduced in the US [67].

658 The result of applying the method to both datasets is a time series containing only gas usage matching the
659 identified pattern of space heating usage, and free from any anomalous spikes. The effect of these events on
660 the internal temperature matches predictions based on an optimised dynamic model to a high level,
661 indicating that the results are likely accurate. An experimental validation exercise was also carried out with
662 positive results. Though this may not be sufficient to measure the performance of our implementation in a
663 precise fashion, it does indicate that this method, or ones based on it, are a successful approach to isolating
664 heating gas usage without resorting to any additional data, which makes the use of smart meter data for a
665 dynamic model of energy performance possible at high frequencies with a significantly reduced uncertainty,
666 and effectively removes the need for the cumbersome instrumentation setups that have been required so
667 far, in cases where meter readings are sufficiently accurate.

668 Regarding general applicability, the observed characteristics of the heat series seem to stem mostly from the
669 fact that the boiler generating it is a combi boiler, which may give cause to limit the intended domain of
670 application to this type of boilers, but the method is general enough to be adaptable to different types, as
671 long as their behaviour while heating hot water (likely in a cylinder) is distinguishable from that of space

672 heating. In the available data, we see that although hot water heating did occur in the form of spikes most of
673 the time, longer loads (baths) produce different patterns. Additionally, hot water spikes were effectively
674 identical to those observed as parts of the heating cycle, which points to a strength of this method; its
675 reliance on dynamic variation prevents otherwise confounding behaviours from causing false detection.
676 Nonetheless, it is still necessary to test on different data sources to validate and improve it, as several
677 factors may be less than ideal in other setups. The window detection method, in particular, though seeming
678 reasonable for any signature, may be more error-prone with different forms of heating control or with an
679 increased number of gas-powered systems. The nature of operation in House 2 is already challenging to
680 capture with without careful tuning, but a more sophisticated filter would likely manage to resolve this
681 issue. This is not likely to have an impact on the applicability of the method in general, as heating events are
682 expected to be the most pronounced instances of sustained use in any gas usage time series; but for the
683 sake of designing an automatic optimisation scheme that is applicable in all cases, a method with more
684 configurable parameters may yield more precise results. Though we have suggested several optimisation
685 options for automating parameter selection, this is not a problem that can generally be solved confidently
686 with such a limited dataset.

687
688 Nonetheless, given the low number of configurable operations, an automated robust implementation seems
689 within reach. This can be realised through possible future development of this method to include automatic
690 detection of signatures based on clustering, which has been demonstrated in various applications before,
691 including disaggregation and energy signatures [68, 69, 70]. It is possible that a set of simple heuristics aided
692 by a clustering algorithm could find heating signatures automatically and reliably, which could in turn prove
693 important for datasets captured at a lower sampling frequency where patterns would be less well-defined,
694 resulting in a reduced requirement for sampling frequency.

695
696 Since reliability of disaggregation seems to be good, there is indeed potential for gas disaggregation to
697 become part of the official smart meter infrastructure. In practice, a universal deployment might take the
698 form of a software component running locally on smart-meter processors or on existing home energy
699 monitors, in order to make use of 1-minute (or less) frequency data without burdening the network with a
700 large amount of readings.

701 **Conclusion**

702
703 A method for disaggregating space heating demand from total gas meter readings is presented, developed
704 and tested using data from a case study in two houses in the south west of England. It is found that the
705 profile of use for the installed gas-fired heating systems is distinct and consistent; easily identified in the
706 complete time series. Automatically detecting all possible windows of heating activity and comparing them
707 to an identified signature was found to produce credible results, supported by the available evaluation
708 methods. Though the method is intended to be quite generic, further study is recommended with different
709 types of boilers and heating systems in various contexts, to ensure robustness and enable the creation of an
710 automated optimisation routine to automatically detect heating signatures and tune selection and filtering
711 parameters.

712
713 The success of the presented method in achieving the required separation between the gas load of the
714 heating system and that of other appliances could allow studies of building performance evaluation using
715 smart-meter data in houses where heating is provided by gas. Analysis using dynamic models with this data
716 is possible, should the other required inputs be available, and is expected to be of much higher use for
717 diagnosing performance gaps than typical steady state ones.

718 719 **Acknowledgment**

720 This work has been generously funded by the British Council Newton Fund, as part of the Institutional Links
721 program, award reference 172732496.

722 723 724 **Declaration of interest**

725 The authors have no financial or personal relationships with any individual or organisation, or any other
726 interests that could inappropriately influence this work. The funding body financing this work had no
727 influence on its design, execution, or decisions relating to publishing.

728 **References**

729

- [1] M. Economidou, B. Atanasiu, C. Despret, J. Maio, I. Nolte, O. Rapf and others, "Europe's buildings under the microscope. A country-by-country review of the energy performance of buildings," *Buildings Performance Institute Europe (BPIE)*, 2011.
- [2] M. Maivel, K. Kuusk, R. Simson, J. Kurnitski, T. Kalamees and others, "Status on the Ground-Overview of existing surveys on energy performance related quality and compliance," *QUALICheck Project Report*, 2015.
- [3] P. De Wilde, "The gap between predicted and measured energy performance of buildings: A framework for investigation," *Automation in Construction*, vol. 41, pp. 40-49, 2014.
- [4] E. Mangematin, G. Pandraud and D. Roux, "Quick measurements of energy efficiency of buildings," *Comptes Rendus Physique*, vol. 13, pp. 383-390, 2012.
- [5] P. Boisson and R. Bouchié, "ISABELE method: In-Situ Assessment of the Building Envelope performances," in *9th International Conference on System Simulation in Buildings*, 2014.
- [6] J. K. Kissock, J. S. Haberl and D. E. Claridge, "Inverse modeling toolkit: numerical algorithms (RP-1050)," *Transactions-American society of heating refrigerating and air conditioning engineers*, vol. 109, pp. 425-434, 2003.
- [7] M. F. Fels, "PRISM: an introduction," *Energy and Buildings*, vol. 9, pp. 5-18, 1986.
- [8] R. Jack, Building diagnostics: practical measurement of the fabric thermal performance of houses, PhD Thesis, PhD, University of Loughborough, 2015.
- [9] J. Berger, S. T. Tascia-Guernouti and M. Humbert, "Experimental method to determine the energy envelope performance of buildings," 2010.
- [10] D. Butler and A. Dengel, Review of Co-heating Test Methodologies: Primary Research, NHBC Foundation, 2013.
- [11] S. Hammarsten, "A critical appraisal of energy-signature models," *Applied Energy*, vol. 26, pp. 97-110, 1987.
- [12] R. Everett, "Rapid thermal calibration of houses report ERG055," *Science and Engineering Research Council Report. London*, 1985.
- [13] C. P. Underwood and F. W. H. Yik, "Modelling Heat Transfer in Building Envelopes," in *Modelling Methods for Energy in Buildings*, Blackwell Publishing Ltd, 2008, pp. 47-92.

- [14] G. Bauwens, *In Situ Testing of a Building's Overall Heat Loss Coefficient-Embedding Quasi-stationary and Dynamic Tests in a Building Physical and Statistical Framework*, 2015.
- [15] J. J. Bloem, *System identification applied to building performance data*, Office for Official Publications of the European Communities, 1994.
- [16] R. R. Crawford and J. E. Woods, "A method for deriving a dynamic system model from actual building performance data," *ASHRAE Trans.:(United States)*, vol. 91, 1985.
- [17] A. Afram and F. Janabi-Sharifi, "Theory and applications of HVAC control systems – A review of model predictive control (MPC)," *Building and Environment*, vol. 72, pp. 343-355, 2014.
- [18] M. D. Knudsen and S. Petersen, "Demand response potential of model predictive control of space heating based on price and carbon dioxide intensity signals," *Energy and Buildings*, vol. 125, pp. 196-204, 2016.
- [19] H. Madsen, P. Bacher, G. Bauwens, A.-H. Deconinck, G. Reynders, S. Roels, E. Himpe and G. Leth, "Thermal Performance Characterization using Time Series Data-IEA EBC Annex 58 Guidelines," International Energy Agency, 2015.
- [20] H. Harb, N. Boyanov, L. Hernandez, R. Streblow and D. Müller, "Development and validation of grey-box models for forecasting the thermal response of occupied buildings," *Energy and Buildings*, vol. 117, pp. 199-207, 2016.
- [21] P. A. Strachan and L. Vandaele, "Case studies of outdoor testing and analysis of building components," *Building and Environment*, vol. 43, pp. 129-142, 2008.
- [22] P. R. Armstrong, J. A. Dirks, R. W. Reilly, J. W. Currie, R. J. Nesse, B. Nekrasov and O. V. Komarov, "Russian apartment building thermal response models for retrofit selection and verification," *ACEEE Summer Study on Energy Efficiency in Buildings, Monterey, CA*, 2000.
- [23] V. Dimitriou, S. K. Firth, T. M. Hassan, T. Kane and F. Fouchal, *Developing suitable thermal models for domestic buildings with Smart Home equipment*, The Bartlett, UCL Faculty of the Built Environment Institute for Environmental Design and Engineering (© IBPSA-England), 2014.
- [24] W. Swan, R. Fitton and P. Brown, "A UK practitioner view of domestic energy performance measurement," *Proceedings of the Institution of Civil Engineers - Engineering Sustainability*, vol. 168, pp. 140-147, 2015.
- [25] D. Connolly, "Heat Roadmap Europe: Quantitative comparison between the electricity, heating, and cooling sectors for different European countries," *Energy*, vol. 139, pp. 580-593, 2017.
- [26] S. R. Vitullo, *Disaggregating time series data for energy consumption by aggregate and individual customer*, PhD Thesis, Marquette University, 2011.
- [27] M. Wytock and J. Z. Kolter, "Contextually Supervised Source Separation with Application to Energy Disaggregation.," in *AAAI*, 2014.
- [28] A. Zoha, A. Gluhak, M. A. Imran and S. Rajasegarar, "Non-intrusive load monitoring approaches for

disaggregated energy sensing: A survey,” *Sensors*, vol. 12, pp. 16838-16866, 2012.

- [29] S. R. Vitullo, R. H. Brown, G. F. Corliss and B. M. Marx, “Mathematical models for natural gas forecasting,” *Canadian applied mathematics quarterly*, vol. 17, pp. 807-827, 2009.
- [30] S. Yamagami, H. Nakamura and A. Meier, “Non-Intrusive Submetering of Residential Gas Appliances,” in *Proceedings of the ACEEE Summer Study on Energy Efficiency in Buildings*, 1996.
- [31] R. Sonderegger, “Diagnostic tests determining the thermal response of a house,” *Lawrence Berkeley National Laboratory*, 2010.
- [32] A. Rabl, “Parameter estimation in buildings: Methods for dynamic analysis of measured energy use,” *J. Sol. Energy Eng.:(United States)*, vol. 110, 1988.
- [33] P. Bacher, P. A. Saint-Aubain, L. E. Christiansen and H. Madsen, “Non-parametric method for separating domestic hot water heating spikes and space heating,” *Energy and Buildings*, vol. 130, pp. 107-112, 2016.
- [34] K. K. Andersen, H. Madsen and L. H. Hansen, “Modelling the heat dynamics of a building using stochastic differential equations,” *Energy and Buildings*, vol. 31, pp. 13-24, 2000.
- [35] M. Killian and M. Kozek, “Ten questions concerning model predictive control for energy efficient buildings,” *Building and Environment*, vol. 105, pp. 403-412, 2016.
- [36] G. Mustafaraj, G. Lowry and J. Chen, “Prediction of room temperature and relative humidity by autoregressive linear and nonlinear neural network models for an open office,” *Energy and Buildings*, vol. 43, pp. 1452-1460, 2011.
- [37] R. BOUCHIE, F. Alzetto, A. Brun, C. Weeks, M. Preece, A. Muhammad and M. Sisinni, “Methodologies for the Assessment of Intrinsic Energy Performance of Buildings Envelope,” March 2015. [Online]. Available: http://performerproject.eu/wp-content/uploads/2015/04/PERFORMER_D1-2_Methodologies-for-the-Assessment-of-Intrinsic-Energy-Performance-of-Buildings-Envelope.pdf. [Accessed 15 May 2018].
- [38] P. Vigo and F. Cascetta, “Sensors for measuring flow,” *Sensors Set: A Comprehensive Survey*, pp. 373-436, 1993.
- [39] G. Ficco, “Metrological performance of diaphragm gas meters in distribution networks,” *Flow Measurement and Instrumentation*, vol. 37, pp. 65-72, 2014.
- [40] G. Buonanno, “On field characterisation of static domestic gas flowmeters,” *Measurement*, vol. 27, pp. 277-285, 2000.
- [41] H. Kershaw, T. Lelyveld, S. Burton, G. Orr, H. Charlick, T. Dennish and M. Crowther, “In-situ monitoring of efficiencies of condensing boilers-TPI control project extension,” *Energy Saving Trust and Department of Energy \& Climate Change*, 2010.
- [42] B. Sodagar and D. Starkey, “The monitored performance of four social houses certified to the Code for Sustainable Homes Level 5,” *Energy and Buildings*, vol. 110, pp. 245-256, 2016.

- [43] A. Beizaee, D. Allinson, K. J. Lomas, E. Foda and D. L. Loveday, "Measuring the potential of zonal space heating controls to reduce energy use in UK homes: The case of un-furbished 1930s dwellings," *Energy and Buildings*, vol. 92, pp. 29-44, 2015.
- [44] M. Zeifman and K. Roth, "Non-Intrusive Appliance Load Monitoring (NIALM): Review and Outlook," *IEEE Transactions on Consumer Electronics*, vol. 75, no. 1, pp. 76-84, 2011.
- [45] S. R. Shaw, S. B. Leeb, L. K. Norford and R. W. Cox, "Nonintrusive Load Monitoring and Diagnostics in Power Systems," *IEEE Transactions on Instrumentation and Measurement*, vol. 57, pp. 1445-1454, 7 2008.
- [46] G. W. Hart, "Nonintrusive appliance load monitoring," *Proceedings of the IEEE*, vol. 80, pp. 1870-1891, 1992.
- [47] V. Stankovic, J. Liao and L. Stankovic, "A graph-based signal processing approach for low-rate energy disaggregation," in *2014 IEEE Symposium on Computational Intelligence for Engineering Solutions (CIES) Proceedings*, 2014, pp. 81-87.
- [48] Kamstrup, "MULTICAL 403: Specifications," Kamstrup, 2017. [Online]. Available: <https://www.kamstrup.com/en-en/products-solutions/thermal-energy-meters/multical-403>. [Accessed 15 May 2018].
- [49] E. Instromet, "BK-G40 Product Specifications," 11 2017. [Online]. Available: https://www.elster-instromet.com/de/product-details/165/de/BK-G40,_BK-G65,_BK-G100. [Accessed 15 May 2018].
- [50] M. Baranski and J. Voss, "Genetic algorithm for pattern detection in NIALM systems," in *Systems, man and cybernetics, 2004 IEEE international conference on*, 2004.
- [51] K. A. Nguyen, R. A. Stewart and H. Zhang, "An Intelligent Pattern Recognition Model to Automate the Categorisation of Residential Water End-use Events," *Environ. Model. Softw.*, vol. 47, pp. 108-127, 9 2013.
- [52] E. Thomaz, V. Bettadapura, G. Reyes, M. Sandesh, G. Schindler, T. Plötz, G. D. Abowd and I. Essa, "Recognizing water-based activities in the home through infrastructure-mediated sensing," in *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, 2012.
- [53] M. Tewolde, J. P. Longtin, S. R. Das and S. Sharma, "Determining appliance energy usage with a high-resolution metering system for residential natural gas meters," *Applied Energy*, vol. 108, pp. 363-372, 2013.
- [54] G. Cohn, S. Gupta, J. Froehlich, E. Larson and S. N. Patel, "GasSense: Appliance-level, single-point sensing of gas activity in the home," in *International Conference on Pervasive Computing*, 2010.
- [55] G. M. Huebner, J. Cooper and K. Jones, "Domestic energy consumption—What role do comfort, habit, and knowledge about the heating system play?," *Energy and Buildings*, vol. 66, pp. 626-636, 2013.
- [56] E. Keogh and S. Kasetty, "On the need for time series data mining benchmarks: a survey and empirical demonstration," *Data Mining and knowledge discovery*, vol. 7, pp. 349-371, 2003.

- [57] D. J. Berndt and J. Clifford, "Using dynamic time warping to find patterns in time series.," in *KDD workshop*, 1994.
- [58] B.-H. Juang, "On the hidden Markov model and dynamic time warping for speech recognition—A unified view," *Bell Labs Technical Journal*, vol. 63, pp. 1213-1243, 1984.
- [59] J. Aach and G. M. Church, "Aligning gene expression time series with time warping algorithms," *Bioinformatics*, vol. 17, pp. 495-508, 2001.
- [60] E. G. Caiani, A. Porta, G. Baselli, M. Turiel, S. Muzzupappa, F. Pieruzzi, C. Crema, A. Malliani and S. Cerutti, "Warped-average template technique to track on a cycle-by-cycle basis the cardiac filling phases on left ventricular volume," in *Computers in Cardiology 1998*, 1998.
- [61] G. Elafoudi, L. Stankovic and V. Stankovic, "Power disaggregation of domestic smart meter readings using dynamic time warping," in *6th International Symposium on Communications, Control and Signal Processing (ISCCSP), 2014*, IEEE, 2014.
- [62] E. Keogh and C. A. Ratanamahatana, "Exact indexing of dynamic time warping," *Knowledge and Information Systems*, vol. 7, pp. 358-386, 01 3 2005.
- [63] V. Chandola, A. Banerjee and V. Kumar, "Anomaly Detection: A Survey," *ACM Comput. Surv.*, vol. 41, pp. 15:1--15:58, 7 2009.
- [64] W. K. Lee, G. S. K. Fung, H. Y. Lam, F. H. Y. Chan and M. Lucente, "Exploration on load signatures," in *International conference on electrical Engineering (ICEE)*, 2004.
- [65] V. Chandola, A. Banerjee and V. Kumar, "Anomaly detection for discrete sequences: A survey," *IEEE Transactions on Knowledge and Data Engineering*, vol. 24, pp. 823-839, 2012.
- [66] M. Pullinger, H. Lovell and J. Webb, "Influencing household energy practices: a critical review of UK smart metering standards and commercial feedback devices," *Technology Analysis & Strategic Management*, vol. 26, pp. 1144-1162, 2014.
- [67] D. S. Sayogo and T. A. Pardo, "Understanding Smart Data Disclosure Policy Success: The Case of Green Button," in *Proceedings of the 14th Annual International Conference on Digital Government Research*, New York, NY, USA, 2013.
- [68] D. Minnen, T. Starner, I. Essa and C. Isbell, "Discovering characteristic actions from on-body sensor data," in *Wearable computers, 2006 10th IEEE international symposium on*, 2006.
- [69] F. Gullo, G. Ponti, A. Tagarelli, S. M. Ruffolo and D. Labate, "Low-voltage Electricity Customer Profiling Based on Load Data Clustering," in *Proceedings of the 2009 International Database Engineering & Applications Symposium*, New York, NY, USA, 2009.
- [70] E. Himpe and A. Janssens, "Data-Driven Modelling of the Energy Use in Dwellings using Smart Meter Data," in *CLIMA 2016-12th REHVA World Congress*, 2016.

ⁱ Defined as the rate of heat loss (W) from the entire envelope of a building per Kelvin of temperature difference between the internal and external environment.