



Measuring the energy intensity of domestic activities from smart meter data



L. Stankovic^{a,*}, V. Stankovic^a, J. Liao^a, C. Wilson^b

^a Department of Electronic and Electrical Engineering, University of Strathclyde, 204 George St., Glasgow G1 1XW, UK

^b Tyndall Centre for Climate Change Research, School of Environmental Sciences, University of East Anglia, Norwich NR4 7TJ, UK

HIGHLIGHTS

- Innovative method linking appliance usage and energy use with domestic activities.
- Inferring the energy and time use profile of activities based on smart meter data.
- Standardised metrics quantifying energy intensity + temporal routines of activities.
- Insights from analysing electricity consumption through the lens of activities.

ARTICLE INFO

Article history:

Received 11 June 2016

Received in revised form 9 September 2016

Accepted 26 September 2016

Keywords:

Activity recognition

Energy disaggregation

Time-use statistics

Energy intensity

Energy efficiency metrics

ABSTRACT

Household electricity consumption can be broken down to appliance end-use through a variety of methods such as modelling, sub-metering, load disaggregation or non-intrusive appliance load monitoring (NILM). We advance and complement this important field of energy research through an innovative methodology that characterises the energy consumption of domestic life by making the linkages between appliance end-use and activities through an ontology built from qualitative data about the household and NILM data. We use activities as a descriptive term for the common ways households spend their time at home. These activities, such as cooking or laundering, are meaningful to households' own lived experience. Thus, besides strictly technical algorithmic approaches for processing quantitative smart meter data, we also draw on social science time use approaches and interview and ethnography data. Our method disaggregates a households total electricity load down to appliance level and provides the start time, duration, and total electricity consumption for each occurrence of appliance usage. We then make inferences about activities occurring in the home by combining these disaggregated data with an ontology that formally specifies the relationships between electricity-using appliances and activities. We also propose two novel standardised metrics to enable easy quantifiable comparison within and across households of the energy intensity and routine of activities of interest. Finally, we demonstrate our results over a sample of ten households with an in-depth analysis of which activities can be inferred with the qualitative and quantitative data available for each household at any time, and the level of accuracy with which each activity can be inferred, unique to each household. This work has important applications from providing meaningful energy feedback to households to comparing the energy efficiency of households' daily activities, and exploring the potential to shift the timing of activities for demand management.

© 2016 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Energy efficiency and energy conservation are priorities for governments worldwide, and have motivated intensive research over

the past decade into understanding how energy is consumed, and how to translate that knowledge into meaningful information to enable energy consumers to take responsibility for their energy consumption. The smart grid concept and smart metering have brought significant advances with respect to integration of heterogeneous and distributed renewable energy sources, and power distribution and energy efficiency in commercial buildings. However, less success has been achieved in the domestic sector [1], despite the fact that this sector accounts for around 30% of total energy

* Corresponding author.

E-mail addresses: lina.stankovic@strath.ac.uk (L. Stankovic), vladimir.stankovic@strath.ac.uk (V. Stankovic), jing.liao@strath.ac.uk (J. Liao), charlie.wilson@uea.ac.uk (C. Wilson).

consumption, for example, 34% in the UK [2], 27% in Spain, 35% in Norway [3], and 37% in USA [4].

Smart meter data provide real-time information on aggregate energy consumption in homes. Disaggregating smart meter data via intrusive or non-intrusive means [5] helps understand how appliances consume electricity in individual households. Recent studies of energy-related feedback have found that electricity consumption data, aggregated or disaggregated down to appliance level, is not often meaningful to households as it is not tied to their lived experience [6,7]. Social scientists have argued that domestic energy use is the largely invisible consequence of deeply embedded social practices occurring within the home [8,9]. Activities such as cooking, washing, listening to music or playing computer games are more consistent with households' own experiences of life at home. Activities are a simple descriptive term for these common ways in which households spend their time [10]. Activities are also used in time-use statistics collected by the national statistical agencies to characterise domestic life [11]. Providing information on energy use through the lens of activities should resonate more clearly with households. Moreover, activities are a more stable constituent of domestic life whereas appliances may be commonly replaced or retrofitted [12].

In this paper, we develop a data sensing and processing methodology, based on collection and analysis of quantitative and qualitative data, to relate electricity consumption to domestic activities. Our methodology makes inferences about electricity-consuming activities following an ontology relating one or more appliances to a particular activity.

An important novelty of our activity-inference method is that it uses disaggregated electricity data from the individual household and so does not rely on stochastic models [13], or on time diaries that are the common method for analysing domestic activities [14]. We apply our method to a sample of 10 homes using data over a one-month period (October 2014). To help compare activity inferences within and between homes, we develop standardised metrics of the energy intensity, timing and duration of domestic activities. These build on proposed metrics of appliance usage intensity [14] but extend their application to regularly occurring domestic activities. Our methodology and results shed light on how routine domestic life constitutes domestic electricity demand and discusses the observed linkages between appliance end-use and activities.

The contributions of this paper can be summarised as follows. We propose an autonomous activity inference algorithm that characterises the energy consumption of domestic life by making the linkages between appliance end-use and activities through an ontology built from qualitative data about the household and disaggregated electricity consumption data. This comprises four key steps, namely capturing quantitative and qualitative data, disaggregating aggregate smart meter measurements via non-intrusive load monitoring where appliance-level measurements are not available, building a household-specific ontology from qualitative data, and finally making inferences, automatically, via an activity recognition algorithm using the ontology and disaggregated appliance end-use data. We identify and characterise sources of uncertainty in making inferences using the proposed methodology, and discuss the limitations of the methodology. We also propose two novel standardised metrics to enable easy quantifiable comparison within and across households of the energy intensity and routine of activities of interest. Finally, we demonstrate our results over a sample of ten households with an in-depth analysis of which activities can be inferred with the qualitative and quantitative data available for each household at any time, and the level of accuracy with which each activity can be inferred, unique to each household.

The paper is organised as follows. Section 2 reviews relevant literature on linking electricity consumption in the home to

appliances and activities. Section 3 details the key steps in our novel activity-inference methodology and linkages between activities and appliance end-use. Section 4 explains its application to a sample of ten homes, the activities which could be inferred, and data limitations. Section 5 introduces standardised metrics for comparing the energy intensity and routine of activities within and between households, before discussing quantifiable insights from analysing electricity consumption through the lens of activities. Section 6 concludes and draws out key implications for further research.

2. Literature review

Understanding domestic electricity demand and consumption is enabled by methods that (1) focus on appliance usage through load disaggregation, or (2) explain electricity usage in terms of domestic activities. We review each body of literature in turn.

2.1. Understanding electricity consumption through load disaggregation

Disaggregation of meter data helps understand the underlying constituents of electricity consumption, rather than just describing it at the aggregate level. Load disaggregation effectively breaks the electricity consumption of a household down to the individual appliances that contribute to the total load at any point in time.

Since monitoring consumption at the plug level (i.e., sub-metering) is becoming affordable, it is increasingly possible to directly measure the power load of individual appliances. However, sub-metering alone can be intrusive and non-scalable, especially if monitoring 60 or more appliances in a home. For example, Dunbabin et al. [15] measures the electricity consumption of loads in up to 250 households across England for periods of between one month and one year. Loads monitored include TV and auxiliary devices on stand-by, fridges and freezers, and overall demand in English households broken down into different categories of appliances (such as cold appliances, lighting, audiovisual, information and communication technologies or 'ICTs'). However, despite monitoring numerous home appliances, Dunbabin et al. [15] shows that on average 20% of electricity use is still unaccounted for due to the large numbers of appliances that could not be monitored.

In contrast to intrusive approaches for energy disaggregation requiring monitoring devices or interventions in the home, non-intrusive approaches use algorithms to infer load profiles from smart meter data. Non-Intrusive Appliance Load Monitoring (NILM) [5,16] disaggregates a household's total electricity consumption down to specific appliances. NILM effectively creates virtual power sensors at each appliance using purely software tools.

In principle, NILM using smart meter data can disaggregate which appliances were used, when they were used, for how long, and how much energy they consumed. The performance or accuracy of NILM is dependent on the smart meter data time-resolution, an up-to-date repository of appliance signatures, algorithmic complexity, and robustness of the algorithm to unknown signatures in a house. Designing accurate NILM algorithms is thus challenging, especially when the sampling rate is low (in the order of seconds or minutes, typical of smart metering devices being deployed worldwide), or when many unknown appliances are present in the house. There has been significant progress over the past decade (see [16] for a review of various approaches to NILM) and there are a few commercial solutions (e.g., developed by Bidgely, ITG, EEme, PlotWatt, Watty, etc.) that can account mostly for large consuming appliances through use of bespoke metering hardware and algorithms. More recently, NILM research is focused on

addressing the challenges of accuracy and complexity, increasingly using only active power measurements, available from large scale deployment of smart meters [17]. For example, Zhao et al. [18] presents a low-complexity, practical algorithm that achieves accuracy of 70% or more for a range of common domestic appliances, and calculates the relative contribution of individual appliances to the total load.

2.2. Understanding electricity consumption through activities

In order to improve load shifting and demand management, we need to understand the temporal flexibility of activities giving rise to energy demand [19]. Some activities may have inflexible schedules, such as cooking, unlike other activities, such as laundering, which may have schedules that are more flexible. To explore the linkages between activities, appliances and associated electricity consumption, Durand-Daubin [14] shows that quantitative questionnaires, activity diaries and measured electricity consumption of related appliances, are generally effective separately in measuring intensity and time of use of appliances studied, but together are not always consistent. That is, intensity and time of use collected from time diaries can be very different from load measurements collected from sub-metering. One reason is that many appliances differ from when they are being actively used (captured by time diaries) compared to when they are on or off (captured by sub-metering). This is especially the case for ICTs, like TV and computers, which can remain on for a prolonged period instead of being switched on and off between discrete periods of active usage.

Widen [13] proposes a stochastic model that estimates the load profiles of households by first generating synthetic activity patterns from sample time diaries, and then mapping these patterns to load demand of the household using appliance-level load measurements. This approach captures the stochastic behaviour of demand, in contrast to standard load profiling which is calculated by averaging the meter readings of a number of households. The households are grouped into householder dwelling types, each representing a large population of similar customers. Ellegard [20] builds on such an approach to develop a visualisation software which displays the relationship between activities and appliance use, using time diaries and a probabilistic model for generating householders' behavioural schedules to predict energy consumption [21]. However, there are limitations to these models, namely: (i) the complexity of detailed modelling of cooking, which can involve numerous appliances, is simplified by a constant load demand, (ii) any residual demand that cannot be attributed to activities (e.g., base load, refrigerator, etc.), is approximated by a fixed load, (iii) the model is trained on multiple households for a couple of days only and thus does not fully capture variability in household routines.

More recently, with the availability of real-time data from a range of sensors, including smart meters, a radically different approach is adopted that infers domestic activities from sensor data. Autonomous activity recognition is recognised as an important enabler of smart home technologies [22]. Activity recognition using real-time smart meter data can yield more accurate insights into households' activities and their consequences for energy consumption.

Activity recognition research emerged to enable assisted living and telemedicine and has since primarily focused on healthcare applications to detect physical activities of people through a range of sensors in the smart home. Domestic activity recognition is based on making an inference by fusing information from heterogeneous sensors and resolving uncertainty due to the stochastic nature of human behaviour and imperfect sensing equipment. To detect a domestic activity, multiple heterogeneous sensors, ranging from numerous switch/pressure sensors to occupancy sensors,

sensors for measuring heartbeat, walking patterns, and different environmental sensors are usually installed in the home. Belley et al. [23] uses NILM together with individual appliance energy sensors to infer activities such as making coffee, with variants such as making coffee and making pancakes, or making coffee and drying hair. Using high sampling rates (~ 60 Hz) of active and reactive power, Clement et al. [24] first analyses patterns of usage of particular appliances, and then identifies activities such as shopping, media, food preparation, telephoning, and hygiene. Note that these studies are in the context of assisted living and use specialised high-rate appliance-level power sensors and other environmental sensors such as occupancy.

It is only very recently that activity recognition research has been applied more generally in buildings. A recent survey of energy saving and user activity recognition [22] shows that research efforts have been focused on predicting occupancy and mobility patterns in buildings, including homes, using a range of probabilistic methodologies and simple sensors, such as motion or acoustic sensor data, in order to enable smarter control of energy in buildings, especially for HVAC systems, AC and plug loads. The kind of activities that are inferred in [22] include sleeping, eating, watching TV, using a coffee maker, working and entertaining. All these approaches use a variety of sensors, but not electricity sensors/meters.

In [25], appliance-level power sensors, motion, temperature, and water usage sensors are all used to infer six daily activities computing, sleeping, cooking, watching TV, taking a shower and grooming. For each activity, appliances are identified that are directly and indirectly associated. For example, computing is directly associated with appliances like computers and printers, and indirectly associated with appliances like desk lamps. Cooking is directly associated with appliances like microwaves, ovens, stoves, and indirectly associated with kitchen lights. Several machine learning methods are tested using time duration and wattage of appliances. Results are presented for five months, for one house. In their follow-up work [26], a suffix tree encoding method and clustering analysis is used to associate sensor measurements (namely, power readings, motion and temperature sensors) to activities and remove outliers, and then train three machine learning algorithms using annotated sensor events to predict energy consumption of six activities. The authors show a good match with ground-truth data collected from two smart home testbeds over a period of two months. Similarly, in [27], activities such as watching video, computing, grooming, cooking, mobility, and cleaning are inferred using only appliance-level power sensors over a period of 24 h for one family occupied house. Note that these studies require continuous data from multiple and heterogeneous sensors and do not discuss linkages between activity recognition and energy consumption.

Related human-computer interaction research has quantified energy services consumed in homes [28] or the energy consumption of specific appliances and devices [29]. In [28], energy services (including some activities) such as entertainment, ICT-related activities, lighting, refrigeration and cooking are inferred only from appliance-level power sensors for 4 student flats over a period of 20 days. Algorithms of [30] infer energy services such as cooking, washing, or heating in 4-student flats over a period of 21 days through a combination of cooker-mounted webcams, appliance-level power sensors for specific appliances, environmental and motion sensors. Regardless of the meter, monitors and sensors used for data gathering, resulting inferences about energy services or appliance usage are commonly interpreted using qualitative interview or video data on household routines and behaviours, or on exceptional events identified in the energy data [28,30].

The above approaches to activity recognition mostly use a combination of sensors including appliance-level electricity sensors.

None of the studies above assessed the impact of accuracy of using NILM for activity recognition. Practical implementations of resulting solutions are hindered by lack of scalability due to the volume of sensors needed, the level of intrusiveness of some sensors such as cameras, and the reality of appliance-specific electricity sensors or appliances being moved around in households resulting in erroneous or invalid measurements. The above limitations can be mitigated by minimising reliance on continuous data from multiple sensors, drawing on data that only needs to be updated occasionally, e.g., video ethnography data, usage patterns, and appliance survey, thus minimising the volume of sensors with continuous measurements, which need to be checked for accuracy and validity before inferring activities.

To address the challenges of scalability, intrusiveness and sensor measurement validation, Liao et al. [31] develops an activity recognition approach using only NILM applied to smart meter active power readings and qualitative data that requires only occasional updating such as appliance surveys. Three activities are inferred, namely cooking and home entertainment. With additional heat and humidity sensors, the same method is extended to infer non-electricity using activities such as showering and gas-hob cooking. Results are presented for one family home over a period of three months.

Wilson et al. [7] further demonstrate the value of drawing on qualitative data from household interviews and video ethnography to develop and validate the activity inference process in combination with NILM and a few appliance-level power sensors. The methodology is demonstrated on two homes to infer reliable time profiles of a range of domestic activities. Stankovic et al. [32] develop this approach further to make inferences about twelve activities in six households over a period of one month. They analyse the time profiles of activities both within and between households, and discuss the electricity consumption associated with the different activities.

This paper substantially advances this body of work by a methodology that draws on quantitative and qualitative data, demonstrating its versatility and scope of application for disaggregating activities further than is conventionally done so they correspond more closely with the activities used in official time-use statistics using data from a larger set of homes, and proposes and demonstrates the use of standardised metrics for quantifying and comparing both the energy intensity and frequency or routine of activities constituting domestic life. Compared to the literature reviewed, this paper also describes in more detail the linkages between appliance end-use and activity energy consumption, as well as activity duration.

3. Methodology

In this section, we present a novel methodology to infer the occurrence and associated electricity consumption of domestic activities using smart meter data. At the core of the methodology lies an activity recognition algorithm that identifies appliance usage events: (i) by directly monitoring appliances via individual appliance monitors (IAM); (ii) by using Non-Intrusive appliance Load Monitoring (NILM); and (iii) by defining activity ontologies using qualitative data from interviews and physical home surveys. In this section, we describe our multi-step methodology, and discuss how our methodology helps address the challenges associated with activity recognition from available data.

Our methodology consists of five steps, which are applied separately for each home analysed.

1. Activity selection: Select activities relevant to a specific home from the full set of 10 activities that characterise domestic life.

- 2. Data collection:** Collect real-time electricity data from aggregate electricity meters and IAMs. Collect data on home and household characteristics including routines and appliance use patterns.
- 3. Load disaggregation:** Disaggregate electricity data using NILM to identify operation of all appliances that are not monitored directly by IAMs.
- 4. Activities ontology:** Formally map relationships between activities and appliances to build an 'activities ontology'.
- 5. Activity inferences:** Use activities ontology and disaggregated electricity data to make inferences about when and for how long activities are occurring, and their electricity consumption consequences.

We discuss each of these methodological steps in more detail below.

3.1. Activity selection

This paper focuses on a set of ten activities that constitute the majority of life at home. Our set of activities is based on the disaggregated double-digit codes used by the UK's Office of National Statistics (ONS) in their time-use research [11,33]. From the full ONS activity list, we excluded two types of activity: (i) activities that do not take place within the home (e.g., travel), or only take place within the home under specific circumstances (e.g., volunteering, sport), (ii) activities that are not clearly associated with energy-using appliances (e.g., sleeping, eating). Then, the remaining ONS double-digit codes were aggregated and structured into a set of 10 activities which are linked to specific energy-using appliances and which describe discrete goings-on within the home. Four activities, namely cooking, washing, laundering and cleaning, categorised as Daily Routines, correspond directly to ONS double-digit codes. One activity, 'socialising', corresponds to all the within-home activities under the ONS single digit (aggregated) code for 'Social Life & Entertainment'. The remaining five activities, namely 'watching TV', 'listening to radio', 'games', 'computing' and 'hobbies', categorised as 'Leisure & Computing' correspond to the within-home activities under the ONS single digit codes for 'Hobbies and Games, and Mass Media'. We kept four ICT-related activities separate ('computing', 'games', 'tv', 'radio') as these are expanding rapidly in terms of associated technologies, time use, and impact on electricity consumption. We captured all other non-ICT leisure activities under hobbies.

Table 1 describes our set of 10 activities, and compares our comprehensive set with those activities whose energy consumption was inferred, wholly or in part, before in discrete studies reviewed earlier.

Our set of activities can be distinguished broadly as daily routines (cooking, washing, laundering, cleaning) or as computing and leisure (watching TV, listening to radio or music, playing computer games, all other computing, hobbies). Socialising is an activity that constitutes daily life but is not directly electricity consuming. However, it can be inferred indirectly from linked activities, e.g., listening to music.

Not all activities will be relevant for all homes. The next data collection step uses different sources of data to identify the subset of activities relevant for a specific home.

3.2. Data collection

We collect a combination of quantitative and qualitative data in each home being analysed.

Quantitative data comprise aggregate and individual appliance active power in Watts (W) sampled every 8 s, similar to the specifications of smart meters being rolled out nationally [17]. The data

Table 1

List of activities that have been inferred from sensor or energy data.

Activity label	Description of activity	Inferred in different studies	Inferred in our study
Cooking	Cooking, preparing food and drink, washing up	[25–27,30,31]	Yes
Washing	Showering, washing, dressing	[25,26,30,31]	Yes
Laundering	Doing laundry		Yes
Cleaning	Cleaning or housework <i>other than laundry or washing up</i>	[27]	Yes
Watching TV	Watching TV, video, film (i.e., any audiovisual)	[22,25,26]	Yes
Listening to radio	Listening to radio, music (i.e., any audio)		Yes
Games	Playing games on console, computer, tablet, smartphone		Yes
Computing	Using computer, tablet, smartphone <i>other than for games</i>	[25–27]	Yes
Hobbies	Doing hobbies, sports		Yes
Socialising	Entertaining, socialising, being with people at home	[22,31]	(Yes)

used in this study are publicly available in [34]. Details of our quantitative data collection platform can be found in [35]. Up to nine IAMs were used in each home. The electrical consumption of the remaining appliances used in the home were all obtained via load disaggregation using NILM [36].

Collected qualitative data comprise: (1) appliance surveys; (2) semi-structured household interviews on activities; (3) video ethnography on technology ownership and usage. The appliance surveys are to help identify unknown signatures obtained during NILM. Details of the qualitative data procedures can be found in [10,37]. The interview and video data are coded (analysed and interpreted) in terms of domestic routines and are used primarily for mapping relationships between activities and technologies for each household.

3.3. Load disaggregation

Information on when appliances were running is measured either through IAMs or inferred from the aggregate readings via NILM. We use a mix of physical sensors (IAMs) and virtual sensors (via NILM) for two reasons. First, monitoring every single appliance in a home via a physical sensor is expensive and unpractical. As a result, we use only up to nine IAMs in each home which keeps acquisition, processing and storage cost and complexity manageable. Second, NILM introduces inference uncertainty, which depends on the accuracy of the NILM algorithm and appliance type; as a result, we do not rely exclusively on NILM. In this paper, we use the approach proposed in [36] based on decision tree (DT), which has the advantages of minimal training, low computational cost, and high performance at low sampling rates using active power data only. The DT-based method first needs training, during which, for each known appliance, we detect the maximum increasing and decreasing edges, from which a DT model is designed. Labelling of signatures detected by the NILM algorithm is dependent on data from the household's appliance survey. Training only needs to be performed every time a new appliance is introduced in

the home. Useful training data could potentially also be generated by appliance time diaries completed by households.

The input to the disaggregation algorithm is the time-stamped active aggregate load as well as the DT model for that household, available from the training step. The output of the disaggregation process is detailed information about each appliance use or event, detected within the chosen period of disaggregation. Specifically, this comprises the time when the appliance was switched on and when it was switched off, the duration of that event, appliance label, average effective power (W) and the total consumption (in kW h) of the appliance during that event. The aim of our disaggregation algorithm is to detect accurately as many events as possible to account for electricity-using appliances, which contribute to the aggregate load at any given point in time.

3.4. Activities ontology

Activities cannot be inferred if they lack any direct or indirect association with electrical appliances. The proposed ontology distinguishes associated appliances, which mark an activity taking place at the same time as another activity.

The output of NILM and data from IAMs provide the list of specific appliances used together with their time of use. This information can be related to particular activities using an activities ontology' specific to each home. An activity ontology maps out all known relationships between activities and the appliances used in those activities. The purpose of the ontology is knowability, that is, to link measurable information on appliances to the set of activities characterising everyday life at home. Mappings between appliances and activities are non-exclusive, i.e., one activity can be mapped to one or more appliances, and vice versa. Any given appliance can *definitely, possibly, or indirectly* indicate that an activity is occurring. These are distinguished in the ontology through three corresponding codes: *marker appliance, auxiliary appliance, and associated activity*.

Table 2

Common marker technologies for each activity being inferred.

Activity label	Common marker appliances	Detection method
Cooking	Kettle, microwave, oven, toaster, dishwasher, cooker	NILM
Washing	Electric shower, hair dryer, hair straightener	NILM
Laundering	Washing machine, tumble dryer	NILM
Cleaning	Vacuum cleaner, steam mop	NILM
Watching TV	Television, DVD player, recorder, set top box	NILM, IAM
Listening to radio	CD player, Hi-fi	IAM
Games	Gaming console	IAM
Computing	Computer, printer, scanner	NILM, IAM
Hobbies	Exercise machine, electric drill, sewing machine	NILM
Socialising	n/a	n/a

Marker appliances are appliances whose use tells us when an activity is definitely occurring. For example, a washing machine is one of the marker appliances for the laundering activity. Table 2 shows a general mapping of common marker appliances for all activities in Table 1 and how these marker appliances are measured quantitatively. NILM indicates non-intrusive appliance load monitoring from aggregate meter measurements, and IAM is appliance-level monitoring. Disaggregation cannot capture the use of devices that are highly mobile or that operate on battery power (either permanently or while not plugged in). As a result, mobile or battery-powered devices are not used as marker technologies in the activities ontology.

Appliances used for several different activities cannot be used unambiguously for making activity inferences. The ontology distinguishes marker appliances from auxiliary appliances. Whereas marker appliances identify when an activity is definitely going on, auxiliary appliances indicate that an activity is possibly going on. For example, a householder could use a PC (marker appliance for the computing activity) for the 'listening to radio activity which is defined as any listening to any audio regardless of the device used.

An associated activity refers to the use of an appliance that is a marker for one activity, which is concurrent with or linked to a second activity. For example, a hi-fi is a marker appliance for listening to radio but might also indicate the 'socialising' activity, which is therefore an associated activity for the hi-fi. Thus, associated activities do not have marker appliances. Conventional distinctions between audio, visual, communication, and computing devices are rapidly collapsing. This increases the difficulty of making inferences about specific types of ICT-related activities. To avoid the risk of inference errors, ICT-based activities could be collapsed into a higher order all ICT-related activities but this is less useful as a descriptive characterisation of domestic life.

Table 3 gives an example of part of an ontology in matrix form (ontologies can also be represented diagrammatically, as in [31]). The rows in Table 3 refer to appliances and the columns to activities. A typical ontology might have over 40 rows, one per identified appliance. Each cell of the matrix shows the mapping of relationships between appliances and activities. Marker appliances are shown as an 'x', auxiliary appliances as a '~', and associated activities as an 'o'. Each appliance row includes summary information on its location and general usage patterns if available

from the qualitative data. Each activity column is traffic-light colour coded: green indicates an activity can definitely be inferred; red indicates that an activity is not inferable from the current data; amber refers to an activity that can possibly be inferred if readings are available from IAMs since relevant appliances cannot be reliably inferred by the NILM algorithm (see below for further details).

3.5. Activity inferences: uncertainty and limitations

The input to the activity inference algorithm comprises the appliance label, when the appliance was switched on, when it was switched off, and estimated electricity consumption obtained from the sensor, i.e., disaggregation via NILM or IAMs as well as the ontology. The disaggregation introduces some uncertainty due to NILM's possible misclassification if two appliances have similar active power signatures or due to IAM sensor malfunctions. Uncertainty is also introduced by the stochastic nature of human behaviour, that is, there are many ways an activity can be performed, for example, certain activities may use different subsets of appliances within the defined ontology at different times; this is a common problem in other domestic activity recognition studies. These two uncertainties are termed disaggregation uncertainty and context uncertainty, respectively. Dempster-Shafer (DS) Theory of evidence (see [31] and references therein) provides a useful solution to make reliable inferences given these uncertainties by combining available evidence. DS can make the distinction between uncertain and unknown information and combine evidence from different sources to reach a consensus with some degree of belief. In our case, our physical and virtual sensors comprise the multiple sources of information with uncertainty.

The activity recognition or inference procedure is based on the temporal relationship between different appliances being used. Based on the start time and end time of all appliance usage events, the algorithm groups marker and auxiliary appliances into one activity event. Disaggregation uncertainty, determined by accuracy of the NILM algorithm, and context uncertainty, determined by the likelihood of each marker and auxiliary appliance associated to an activity, are integrated into the model. See [31] for more details on the activity recognition algorithm.

The output of the activity inference procedure is a set of activities, together with their start times and end times to estimate duration. Together with the disaggregated electrical consumption

Table 3 Example of part of an activity ontology for appliances in two rooms of a home. Note (1): x = marker appliance; ~ = auxiliary appliance; o = associated activity. Note (2): green colour coding = inferable; amber = possibly inferable; red = not inferable (see text for details).

Appliance	Signature availability	Measurability	cooking	washing	laundrying	cleaning	watching TV	Listening to radio	games	computing	hobbies	socialising	Appliance information		Typical usage frequency
TV	IAM	yes					x						living room	fixed	daily
hi-fi	IAM	yes						x				~	living room	fixed	2-3 per week
keyboard	NILM	if plug moved									~		living room	fixed	weekly
laptop		no							x		~		living room	mobile	
washing machine	NILM	yes			X								kitchen	fixed	4 per week
dishwasher	NILM	yes	x										kitchen	fixed	daily
microwave	IAM	yes	x										kitchen	fixed	daily
kettle	IAM	yes	x										kitchen	fixed	daily
oven	NILM	yes	x										kitchen	fixed	
toaster	IAM	yes	x										kitchen	fixed	3 per week

obtained from IAM or NILM, we can determine the electrical load associated with each activity from the temporal associations of appliances deemed by DS to form one activity event.

Activity inference is overall assessed with five degrees of uncertainty:

1. **Non-inferable:** Activities associated with non-detectable appliances which cannot be monitored (e.g., a battery-operated appliance like a portable radio) or with mobile, chargeable devices (e.g., laptops or tablets) which are generally not charging when in use. Non-inferable activities are colour coded red in the ontology.
2. **Possibly inferable:** Activities associated with non-detectable appliances because additional quantitative and/or qualitative data is required for disaggregation. One example is ‘washing’ using gas water heating which could be monitored by additional temperature or humidity sensors, or time diaries. Another example is ‘games’ using devices with very low wattage but which could be monitored by an additional IAM. Possibly inferable activities are colour coded amber in the ontology.
3. **Inferable with uncertainty:** Activities associated with appliances whose signatures have not been verified (e.g., using time diaries or an IAM) but can still be detected via disaggregation, or medium-powered appliances whose signatures can get lost in the aggregate data which can make detection sometimes hit and miss. Activities inferable with uncertainty mostly apply to cooking which is associated with a large range of different appliances used at different times and for which all signatures cannot be verified. Activities associated with auxiliary appliances are also classed as inferable with uncertainty because the auxiliary appliance, unlike the marker appliance, indicates that the activity is possibly going on. For example, for some households (as indicated in the ontology via interviews), a computer can be an auxiliary appliance for the ‘watching TV’ activity, and thus ‘watching TV’ for that household would be inferable with uncertainty if the inference algorithm detected the computer running without the TV.
4. **Partially inferable:** Activities associated with gas as well as electricity consumption or with appliances which cannot be disaggregated due to low loads. An example of the first condition is for the ‘washing’ activity associated with both an electric shower and gas-based domestic hot water. We can detect the electric shower via NILM but we cannot always detect the use of domestic hot water for a shower or bath since we are relying on electrical boiler signatures which are disaggregated with some uncertainty. If temperature or humidity sensors were present in the bathroom, washing would be an activity inferable with certainty (see [31]). An example of partially inferable activities associated with appliances with low loads is the ‘listening to radio’ activity which can be partly detected if there is an IAM attached to at least one of appliances associated with that activity (e.g., one CD player is monitored, but other audio devices are not).
5. **Inferable with certainty:** Activities associated with appliances detected reliably via NILM and/or IAMs. Note that the NILM appliances may incur marginal disaggregation error, but usually no more than 10% [36]. All inferable activities - with certainty, partially, with uncertainty - are colour coded green in the ontology.

Time use (e.g., from time diaries) can be very different from load measurements (e.g., from submetering) due to the variable way in which appliances are used in particular activities [14]. With our methodology, we may under-predict activity time use because some appliances may be off during part of the activity, e.g., when

loading the washing machine during the laundering activity, or preparing meat or vegetables for the oven during the cooking activity. On the other hand, our methodology can also overpredict the duration of an activity if marker or auxiliary appliances are on for prolonged periods beyond the duration of an activity, e.g., a radio or TV left on all day, regardless of whether a householder is actually listening to radio or watching TV. As a result, making inferences solely from appliance usage is not reliable as a sole basis for inferring time use. However, given that the disaggregated loads from specific appliances are known and then inferred to be linked to activities based on the ontologies, the energy intensity of domestic activities can be reliably calculated, bearing in mind the uncertainties discussed above and the limitations of underlying appliance usage measurements.

We also note that many activities may be taking place at the same time, for instance cooking, socialising and listening to music. Our methodology allows for multiple activities to be contemporaneous while assigning specific appliances energy use to specific activities. This is intended to be a way of apportioning energy use to discrete activities, and not to split domestic life up into discrete activities. Conversely, a single appliance may be used simultaneously for more than one activity, as in the case of the radio used for listening to radio and socialising. While having one appliance end-use event, we recognise that there are two domestic activities going on at the same time if the ontology indicates that the radio is a marker technology for listening to radio and socialising is an associated activity to the radio. This would also hold true for the case of using a computer that is used for listening to radio, watching TV and computing, when the ontology indicates that the computer is a marker appliance for computing and an auxiliary appliance for listening to radio and watching TV. Note that the only way to capture what the computer was used for (besides installing intrusive usage-monitoring software) is through qualitative data gathered through household interviews and video ethnography on technology ownership and usage. This means that the associations between computer usage and specific activities is potentially different for each household in our sample and will be reflected in the ontology. In the listening to radio and watching TV case, since the computer is an auxiliary appliance to these two activities, we can only infer that listening to radio and watching TV are possibly going on or as described above Inferable with uncertainty. Alternatively, we could use a more aggregated set of activities, which capture all ICT usage into a single generic computing activity. This is even more relevant as conventional distinctions between audio, visual, communication, and computing devices are rapidly collapsing, and thus this increases the difficulty of making inferences about specific types of ICT-related activities. We do not apportion electricity use to associated activities - this is discussed in more detail in Section 5.1.

More generally, aggregating activities helps avoid issues with allocating appliance use to activities, missing inferences, and distinguishing contemporaneous activities. As an example, we categorised our set of ten activities into Daily Routines, Leisure and Computing and Socialising. Analysis at this level of three aggregated categories of activities would clearly be possible from our data. However, aggregation also loses detail and insight into the energy intensity of different domestic activities. We therefore aim to demonstrate our methodologies versatility and scope of application for disaggregating activities so they correspond more closely with the activities used in official time use statistics.

It is worth noting though that the resolution of the analysis is dependent on the ontology, that is, a more disaggregated analysis is dependent on the detail of the interview and ethnographic data which is used to inform the inference algorithms. For example, gaming is not relevant as a disaggregated ICT-related activity for households which do not play computer games.

We also acknowledge that activities cannot account for all actual energy consumption. Heating and lighting are both energy-intensive services but not activities per se. Heating and lighting-related energy use could be apportioned to activities, such as watching TV, taking place in specific rooms for time periods during which those rooms are lit or heated, or could be accounted for separately. We have taken the latter approach in this paper and determine the energy intensity of electric heating, cold appliances, base load and lighting via the residual as separate energy services. This is to help map the actual energy consumption into activities and non-activity energy-consuming services.

4. Implementing the methodology

We use the methodology described to infer up to ten daily activities in a sample of ten households over 24 h daily cycles for each day of the month of October 2014. We selected this period as it was not during the summer (when households are more likely to be outside) and also not during school holidays or festive periods (when domestic activities in households may follow different routines).

4.1. Household sample

Our sample of ten households was selected from a set of twenty households participating in a smart home field trial subject to various criteria, one of which was to capture socio-demographic variability (particularly with respect to household lifecycle, e.g., young families versus retired couples). The focus of this paper is not to capture socio-demographic variation but illustrate the potential of the methodology to do so. Therefore, we ensured that we have some socio-demographic variation within our sample so we could examine similarities or differences between different household types, but any conclusions with respect to socio-demographic differences are indicative only with our small sample size. Households are labelled with numbers between 1 and 21 to be consistent with the REFIT electricity measurement dataset [34] from which the aggregate and appliance-specific electrical load data was obtained. Table 4 summarises the characteristics of each household analysed.

4.2. Activity-related & other electricity consumption

Besides the standard set of activities for the activity inferences (see Section 3.1), we estimated other electricity consumption not accounted for by activities, distinguishing cold appliances, electrical heating (if applicable) – the primary fuel source of heating in the UK is gas, but about 10% of British households supplement with electrical heating in cold winter months [15], base load, and a residual which includes lighting. Activity-related plus other electricity consumption total the aggregate load of a household. Base load is calculated as the minimum level of electricity demand on the household's supply system over 24 h. The residual is the load

calculated from the difference between measured aggregate load and all the activity-related and other electricity consumption. The value of the residual indicates how much we cannot disaggregate, or indirectly account for the total demand of any given household.

4.3. Appliance detection & activities

For each household, an activities ontology was built to map activities which could be inferred with the appliances associated with each activity. Using House 17 as an example, Table 5 shows which activities can be inferred from the collected data, whether our inference algorithm can measure duration (time use) of the activity as well as electricity use, and which appliances were related to activities or other electricity consumption. The appliance information was obtained through a combination of the appliance survey, qualitative data, and NILM. Not all appliances were reported in the appliance survey because they were either unused during the survey or they were not present in the house at the time of the survey.

Fig. 1 shows the relative number of electrical appliances that can be detected reliably via NILM or directly metered from an IAM relative to the total known measurable electrical appliances in each home. On average, 59% of appliances were detected, ranging from 44% in House 21 to 84% in House 17.

Although most high load appliances can be detected, low power appliances (<20 W) such as electric toothbrush, printer, router, DAB radio get lost' in the aggregate data and account for the percentage of appliances that cannot be detected. Another set of appliances that cannot be detected are gas-based, battery-based and mobile appliances such as gas hobs, smart phones, tablets, radios and digital cameras. Detection is also limited by our signature database which only contains signatures we have been able to label and verify via IAMs or appliance time diaries.

4.4. Inferable activities and uncertainties

Having built a detailed activities ontology based on the appliances we could detect via IAM or NILM, we make activity inferences based on the methodology described in Section 3.5. Table 6 shows which of the full set of ten activities could be detected in each home, together with associated uncertainties. As an example, we could detect watching TV in all households with high certainty (coded - 4) because of marker appliances (e.g., TVs) monitored by plug monitors. Similarly, we could detect laundering in all households with high certainty because marker appliances (e.g., washing machines, dryers) have well defined signatures for NILM.

Some activities were inferable through their associations with a different activity (e.g., socialising was inferred indirectly from listening to radio if the two were associated in the ontology for that household). Both activities are recognised as taking place at the same time, however, we assign the electricity use to the listening to radio activity because the radio is the marker technology for this particular activity. We do not assign electricity use to associated activities.

We note that our methodology proposes an activities focus on energy-using domestic life as a complement to, not a substitute for, appliance-level disaggregation and feedback. Arguments from sociological analyses of domestic life and time-use analysis suggest that activities over daily, weekly or seasonal cycles are strongly linked to how households recall or experience their lives at home. Appliance usage, in comparison, is more relevant for households to understand discrete events or moments, or their practices in using a specific appliance. Activity disaggregation is thus intended to complement appliance disaggregation since activities are how households recall or experience their lives at home.

Table 4

Composition of the ten households in our study. Numerical labels for each correspond with labels in publicly-available dataset (see text for details).

House 2	Family of four with two young children
House 3	Two pensioners
House 4	Two pensioners
House 5	Family of four with two children in early teens
House 8	Two pensioners
House 10	Family of four with two young children
House 17	Family of four with one teenage child
House 19	Family of four with two children in early teens
House 20	Family of three adults
House 21	Family of four with two children in early teens

Table 5

List of inferences that can be made (time use and electricity use) about House 17 from the appliances detected, with uncertainty coded from 0 (non-inferable) to 4 (inferable with certainty) - see Section 3.5 for details.

		Inferences			Appliances
		Time use	Electricity use	Uncertainty	
Activity-related	cooking	yes	yes	2	kettle, microwave, coffee maker, toaster, blender, electric cooker, sandwich maker
	washing	yes	yes	2	electric shower, hair straightener
	laundrying	yes	yes	4	washing machine, tumble dryer
	cleaning	no	no	1	n/a
	watching TV	yes	yes	4	television, DVD, set-top box, speakers
	listening to radio	no	no	1	n/a
	games	no	no	0	n/a
	computing	yes	yes	3	desktop computer
	hobbies	no	no	0	n/a
socialising	no	n/a	0	n/a	
Other	cold appliances	n/a	yes		1 fridge-freezer, 2 freezers
	electric heating	n/a	yes		electric heater
	base load	n/a	yes		n/a
	Residual	n/a	yes		n/a

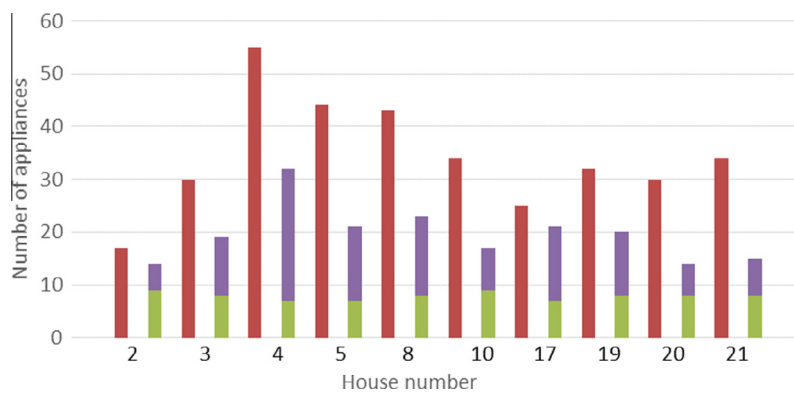


Fig. 1. Number of electrical appliances that can be detected from smart meters in each home for activity inferences: left (red) column shows number of known appliances in the home, right column shows number of appliances detected by IAMS (green) and NILM (purple). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 6

Associated uncertainty level of inferences per activity for each house. Legend: unshaded 0 = non-inferable; quarter-shaded 1 = possibly inferable; half-shaded 2 = inferable with uncertainty; three-quarters shaded 3 = partially inferable; fully shaded 4 = inferable with certainty (see Section 3.5 for details).

House	2	3	4	5	8	10	17	19	20	21
cooking	● 4	● 2	● 2	● 2	● 2	● 2	● 2	● 4	● 3	● 2
washing	● 4	● 3	● 1	● 1	● 4	● 3	● 2	● 1	● 1	● 1
laundrying	● 4	● 4	● 4	● 4	● 4	● 4	● 4	● 4	● 4	● 4
cleaning	● 1	● 2	● 1	● 2	● 2	● 1	● 1	● 1	● 2	● 1
socialising	● 4	● 1	● 1	○ 1	○ 0	○ 0	○ 0	○ 3	○ 0	○ 0
watching TV	● 4	● 4	● 4	● 4	● 4	● 4	● 4	● 4	● 4	● 4
listening to radio	● 4	○ 0	● 1	● 1	● 1	● 1	● 1	● 3	● 1	● 1
ICT-related games	● 1	● 1	○ 0	○ 0	○ 0	○ 0	○ 0	● 4	○ 0	○ 0
computing	● 1	● 1	● 4	● 4	● 4	○ 0	● 3	● 1	● 3	○ 0
hobbies	● 1	● 1	● 2	● 2	○ 0	○ 0	○ 0	● 1	● 1	○ 0

5. Results: the energy intensity of domestic activities

Using our proposed methodology, we generate the time profile of electricity demand for three daily activities and one leisure activity in the ten households in our sample, broken down into

weekday and weekend profiles as illustrated in Fig. 2. There were 23 weekdays and 8 weekends in October 2014.

Cooking occurs throughout the day across all households, but shows clear peak hours for breakfast, lunch and particularly dinner during weekdays and is less structured during weekends across

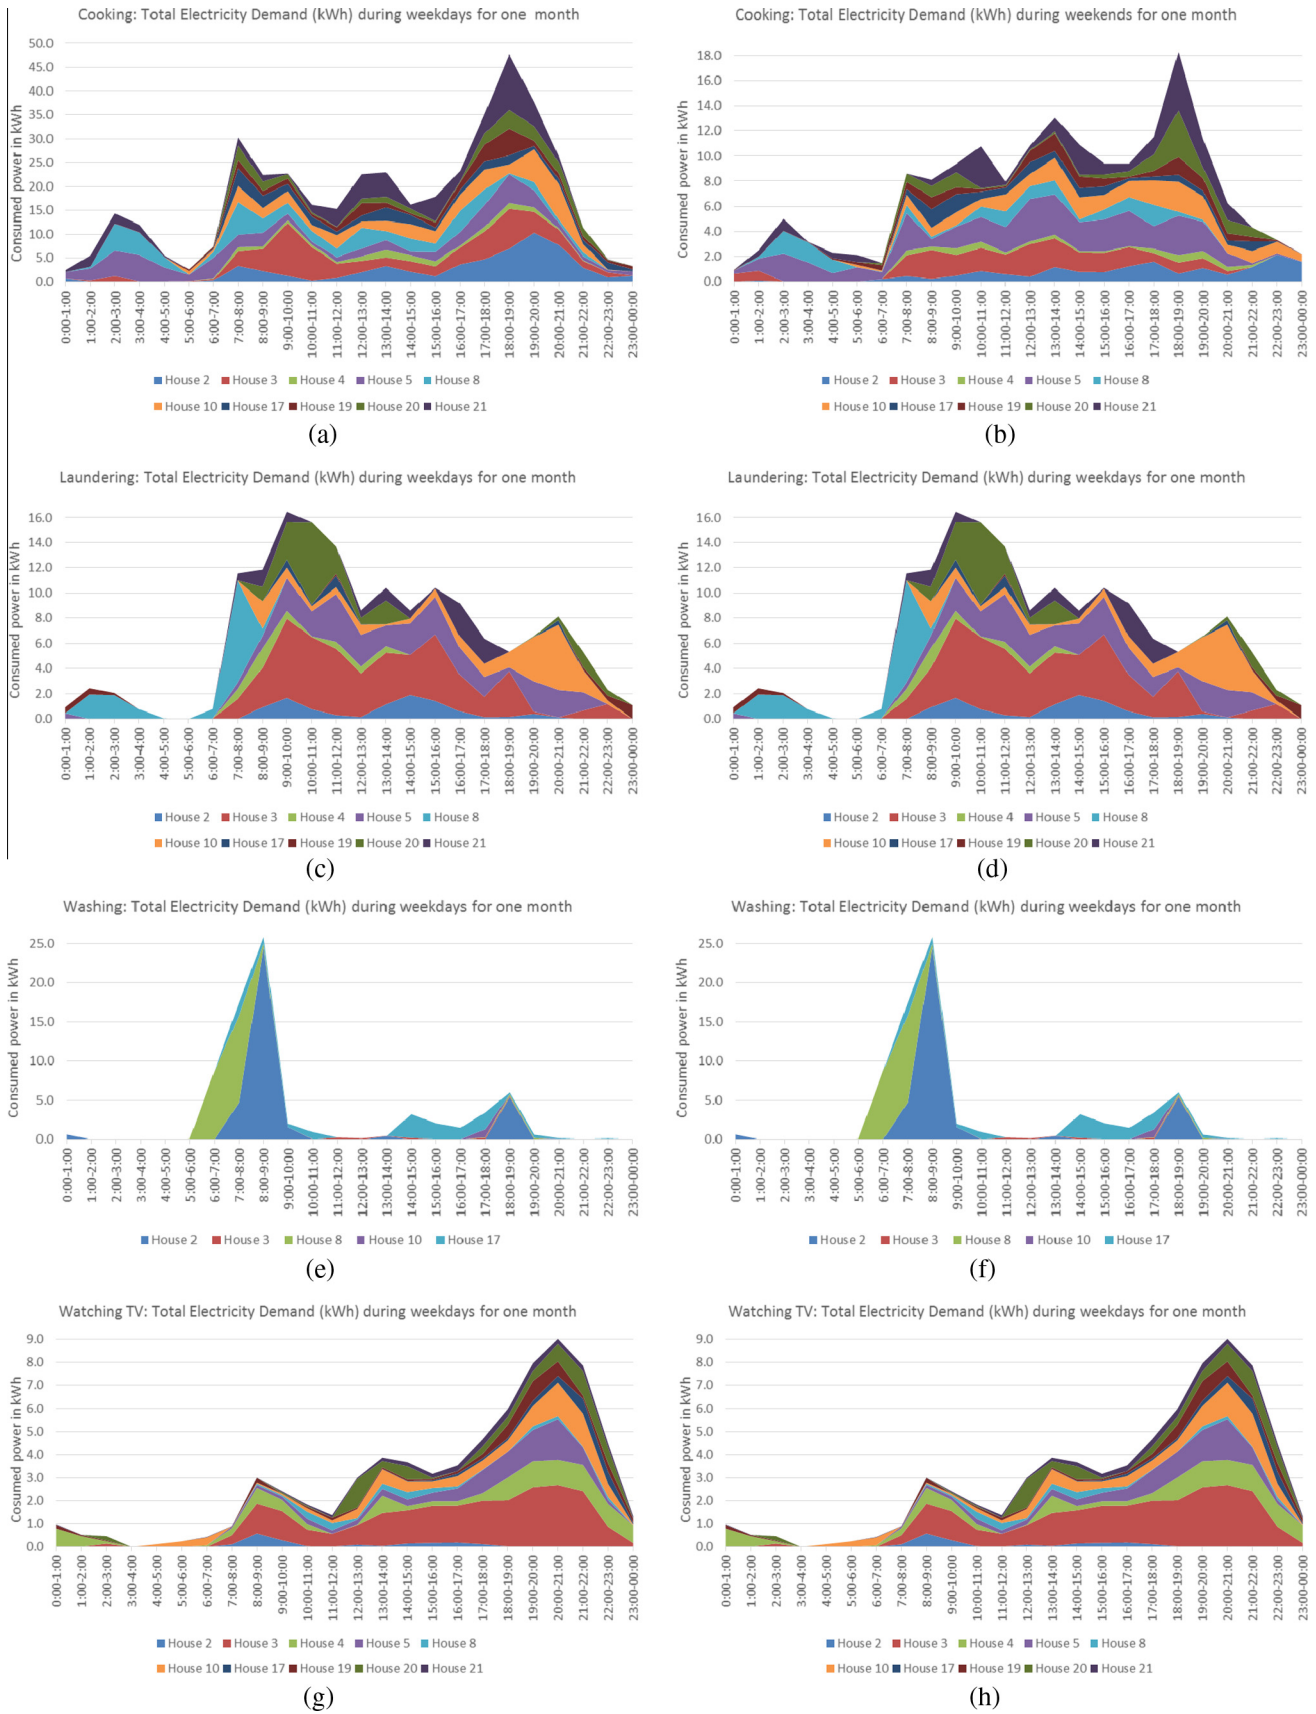


Fig. 2. Hourly activity demand profile of the monthly weekday (left column) and weekend (right column) electricity consumption of cooking, laundering, washing and watching TV. Demand was calculated for all 10 households for all activities shown, except washing which was calculated for 5 households.

most households. Interestingly, watching TV has a similar weekday and weekend demand profile to cooking with a concentration in the evening. Washing has a noticeable peak in the mornings, has clear peaks during weekdays and is less structured during weekends. However, laundering tends to occur primarily mid-morning but is more varied through the day across households unlike the other three daily activities. There is no clear difference in weekday and weekend demand profile, except that peak demand tends to shift more towards later afternoon.

While useful for visually assessing demand and potential for load shifting, Fig. 2 is not immediately helpful for understanding and comparing relative energy consumption within and across households. To demonstrate how our methodology enables standardised quantifiable comparisons in energy use within and across households, we analyse activities using two metrics. First, we define an energy intensity metric to quantify the relative contribution of activities to the total household energy consumption. Motivated by the clear routine observed in households in another study focusing on patterns of use of the kettle [38], we define a routine metric as an indicator of how consistently activities occur during any time period, capturing variability in each activity's duration, frequency given a given time period (i.e., how often the activity occurs) and consistency (i.e., does the activity occur at the same time every day). We demonstrate these metrics using data for a single home in our sample, then show how the metrics enable standardised comparisons of the energy intensity and routine of activities within and between households.

5.1. Metrics of energy-using activities

A measure of *energy intensity* represents the energy consumed by a particular activity normalised to the total aggregate load for a household. We refer to this as energy intensity El_a of Activity a , which is calculated as the percentage of electricity consumed by Activity a during time period T , $E_a(T)$, with respect to the total energy consumption $E(T)$ of the household during that time:

$$El_a(T) = \frac{E_a(T)}{E(T)} 100[\%]. \quad (1)$$

Note that $El_a(T)$ is a real number between 0 and 100, where $El_a(T) = 0$ indicates that Activity a did not occur during time interval T , whereas $El_a(T) = 100\%$ means that the entire household's electricity demand over time period T is attributable to Activity a . For associated activities, we do not calculate energy intensity since the electricity use is already assigned to the activity associated with its marker and auxiliary appliances.

The energy intensity metric allows energy-using activities to be compared consistently between time periods within a household, or within time periods between households. However, it does not take into account the frequency (e.g., how many times in a month during a particular hourly slot) or consistency (e.g., occurs almost every day in a month but at a different time each day) with which activities occur in a time period. Activities may have similar values of $El_a(T)$ if they occur daily for short periods, or only a few days during the time period T but for longer periods. For example, households may launder a few days in a month, whereas cooking generally occurs at similar times every day. To measure the routine occurrence of Activity a within timeslot t over time period T , we use the coefficient of variation, $R_a(t, T)$, also known as the relative standard deviation. We calculate $R_a(t, T)$ as the ratio of the standard deviation of the energy consumed, $\sigma_{E_a(T)}(t) \geq 0$ for Activity a during time period T for each timeslot t to the mean of the energy consumed, $\mu_{E_a(T)}(t) \geq 0$ by Activity a during time period T for each timeslot t :

$$R_a(t, T) = \frac{\sigma_{E_a(T)}(t)}{\mu_{E_a(T)}(t)}. \quad (2)$$

When energy-intensive Activity a does not occur during T , $\mu_{E_a(T)}(t) = 0$ and $\sigma_{E_a(T)}(t) = 0$, thus $R_a(t, T)$ is undefined. Otherwise, $R_a(t, T)$ is a real number ≥ 0 , where $R_a(t, T) = 0$ means that Activity a is occurring always at the same time, every day and consuming the same amount of energy during T . Larger values of $R_a(t, T)$ indicate that an activity occurs less frequently and/or occurs irregularly during timeslot t for a given time period T . Smaller values of $R_a(t, T)$ indicate that an activity occurs frequently with similar durations during timeslot t for a given time period T . In order to calculate $R_a(t, T)$ for each hourly slot (e.g., $t = 10:00$ to $11:00$ AM) during one month (similar to the hourly demand profile in Fig. 2), $\mu_{E_a(T)}(t)$ is calculated by averaging over the whole month, i.e., $\mu_{E_a(T)}(t) = \sum_{i=1}^T \frac{E_a(t_i)}{T}$, where $E_a(t_i)$ denotes the total energy consumed by Activity a during the fixed timeslot t for the i -th day of the month, and $T = 31$ (days) for October 2014. The same principle is applied for calculating $\sigma_{E_a(T)}(t)$.

5.2. Activities within a household

We apply the proposed two metrics of energy intensity and routine to analyse the time profile of energy consumption of particular activities within a single household. We use data from House 5 during October 2014. House 5 is a four-person household with two adults and two children in their early teens. We focus on electricity as set out in our methodology, but in principle, the metrics could also include gas if real-time data were available. We first determine the relative contribution of all inferable activities to the total monthly household consumption. In House 5, the energy intensity for all activities obtained by summing El_a for all a , is 40%. In other words, activities account for 40% of the total monthly electricity load, with cooking and laundering playing a significant part as shown in Fig. 3. Over the whole month, the residual load (including lighting) unaccounted for by activities or other electricity consumption is 28%.

While Fig. 3 is useful to compare the relative energy intensities of different activities over a month, it does not show their distribution over time. Fig. 4 shows the time profile of total monthly electricity consumption for four activities across each hourly time period during a day. Note that the radial axes (Wh) vary for each plot. These rose plots clearly indicate peaks during particular time periods. For example, whereas cooking occurs throughout the day, there is a clear evening peak between 6 and 7 PM. The distinct overnight dishwashing is shown as another peak between 2 and 4 AM (with dishwasher usage assigned to the cooking activity). Laundering is more spread out across the day, whereas hobbies, namely the use of the treadmill, is limited to the early mornings. Watching TV is mainly an evening activity.

Fig. 4 shows how total electricity consumption is apportioned between hourly time slots over a whole month, but it does not capture the routine occurrence of activities $R_a(t, T)$ during any given hourly time period. This is shown in Fig. 5, where $R_a(t, T)$ is averaged over $T = 31$ days for each hourly slot t in a day. Very low values of average $R_a(t, T)$ indicate frequent and consistent occurrences of Activity a . Very high values of $R_a(t, T)$ indicate infrequent and/or inconsistent occurrences of Activity a . Gaps in $R_a(t, T)$ implies that the activity never takes place during that particular hourly slot, e.g., watching TV between midnight and 6 AM.

$R_a(T)$, the 24-h average of $R_a(t, T)$, can be calculated as:

$$R_a(T) = \sum_{t=0:00-1:00}^{23:00-00:00} \frac{R_a(t, T)}{24}. \quad (3)$$

Computing has an almost zero value of $R_{computing}(T) = 0.15$ because the desktop computer is switched on all day everyday, demonstrating consistency and frequency of this activity during a

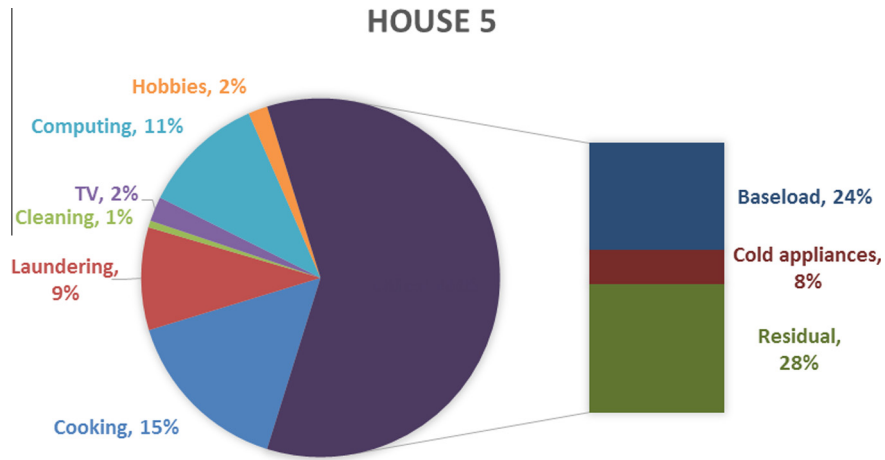


Fig. 3. Breakdown of total monthly electricity load, showing energy intensity $E_i(T)$ of six activities for $T = 31$ days and other electricity consumption (House 5).

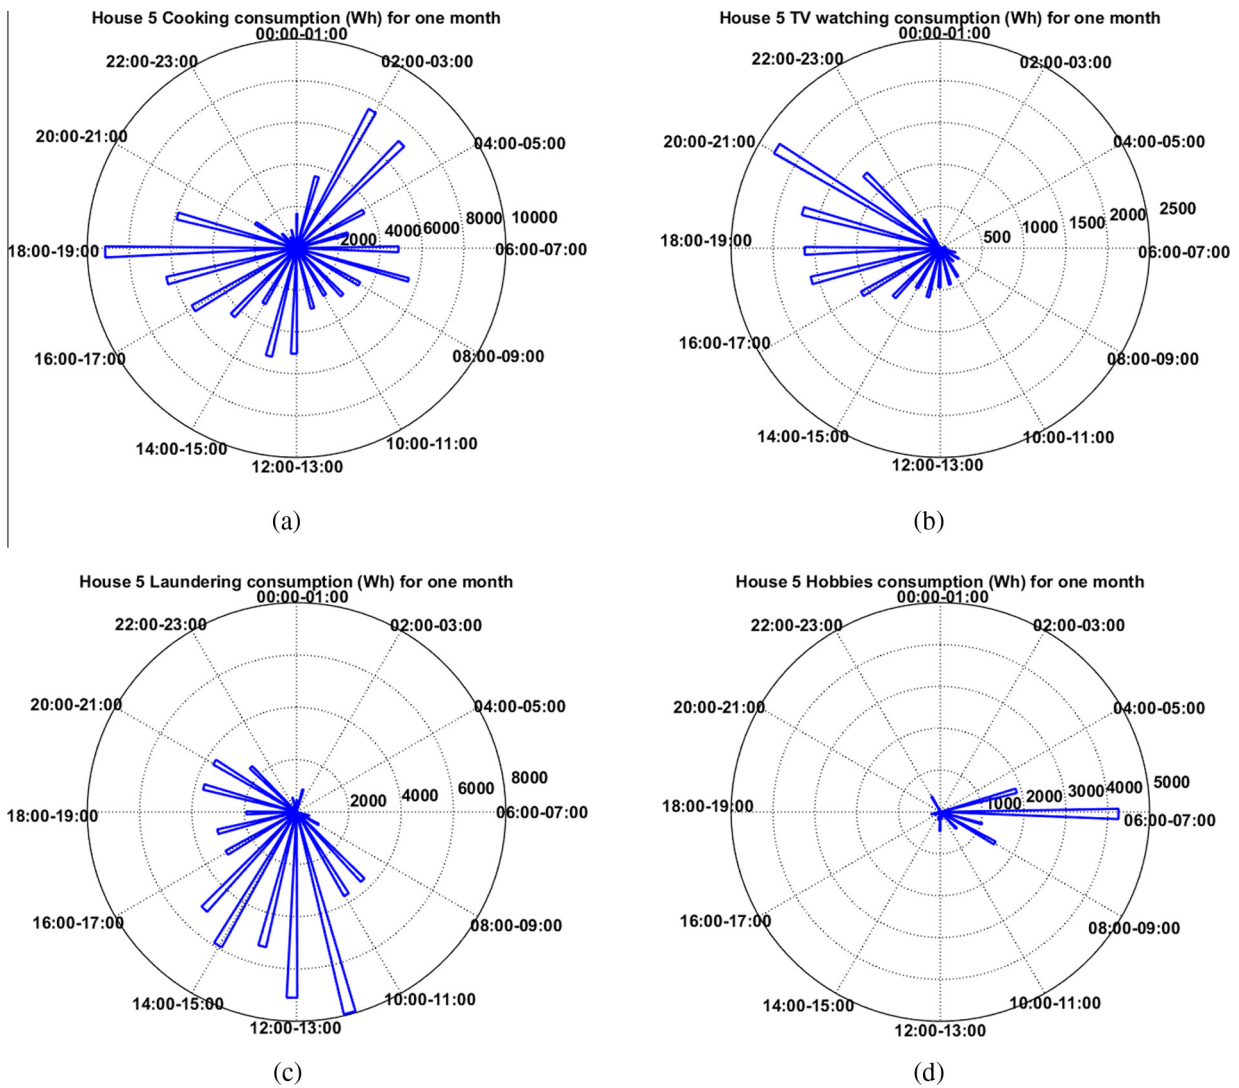


Fig. 4. The time distribution of electricity consumption per activity. Each bin indicates $E_a(T)$ the total electricity consumption (Wh) for $T = 31$ days for each hourly time period (House 5).

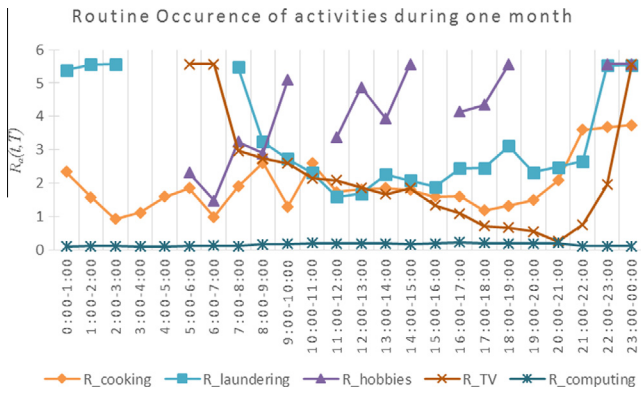


Fig. 5. The hourly routine occurrence of activities, $R_a(t, T)$ over $T = 31$ days (House 5) for hourly timeslots t (x-axis). Note: high values of $R_a(t, T)$ indicate no routine during that time period; lower values indicate routine; gap in values indicates no activities occurring.

month. Watching TV occurs consistently and every day in the evenings with $R_{watchingTV}(T) = 1.75$, and inconsistently at other times of the day. Cooking with $R_{cooking}(T) = 1.93$ happens consistently every day during breakfast and dinner times, and overnight (dishwashing), but is inconsistent at other times of the day. Laundering does not happen every day (infrequent) and also occurs at different times during the day (inconsistent), resulting in a high $R_{laundrying}(T) = 2.77$. Hobbies is consistent between 6 and 8 am, but does not occur every day (infrequent). Hobbies is also inconsistent at other times of the day. The combination of infrequency and inconsistency results in a high $R_{hobbies}(T) = 2.42$, however, $R_{hobbies}(T)$ is less than $R_{laundrying}(T)$ because of the consistent morning routine.

5.3. Activities in ten households

Our activity inference methodology can be used to link electricity consumption to a common set of activities across multiple households. Table 7 shows the energy intensity of all activities inferred (see Table 1) across our sample of ten households. The energy intensity, $EI_a(T)$ for $T = 31$ days, of inferable activities range from 13% to 41% across households. Of all the activities which are generally inferable across households from available electricity data, cooking has the highest energy intensity with an average $EI_{cooking}(T) = 16\%$. Laundering and washing are the next most energy intensive activities. Note that washing could only be inferred in a subset of households due to the use of gas for hot water in other households (see Table 6). Other (non-activity) electricity consumption comprises cold appliances, base load, electrical heating, and a residual, including lighting and charging of portable devices or low-powered devices, which we cannot disaggregate. This residual is 18 to 48% of total electricity use. Lighting in the UK uses an average of 16% of a household's total consumption [39].

In order to understand variability across households for each activity, Fig. 6 shows the $R_a(t, T)$ for $T = 31$ days for each hourly slot t in a day for a sub-sample of households, randomly selected to represent all household compositions (Table 4). The plots have the same interpretation as Fig. 5, which shows various activities for a single household, but Fig. 6 shows a single activity across various households. Small $R_a(t, T)$ values indicate more routine occurrence; large $R_a(t, T)$ indicates less routine occurrence. Convergent or clustered $R_a(t, T)$ values indicate consistency across households; divergent $R_a(t, T)$ indicates variability across households. Cooking generally has a small $R_a(t, T)$ for breakfast, lunch and dinner timeslots, for all households. Houses 5 and 19 watch TV with increasing consistency from early afternoon to evening. Houses 8 and 19 distinctively do laundrying overnight and never after 8 AM. These two

Table 7

EI_a ($T = 31$ days) of inferable activities and contribution of other electricity consumption to total monthly load. Cells with no entry indicate that we had no quantitative and/or qualitative data to make inferences.

a \ House	2	3	4	5	8	10	17	19	20	21
cooking	21%	20%	6%	15%	16%	17%	12%	15%	13%	23%
washing	14%	<1%			6%	<1%	7%			
laundrying	4%	12%	3%	9%	4%	6%	2%	1%	6%	9%
cleaning		1%		1%	1%				1%	
watching TV	1%	7%	2%	2%	1%	3%	1%	2%	3%	1%
listening to radio	<1%							<1%		
ICT-related games								1%		
computing			2%	11%	2%		5%		1%	
hobbies			1%	2%						
% of total electricity use explained by activity inferences	40%	41%	13%	40%	30%	26%	26%	20%	24%	34%
base load	17%	18%	22%	24%	15%	20%	21%	41%	30%	28%
cold appliances	9%	18%	31%	8%	6%	9%	22%	16%	23%	9%
electrical heater					1%	3%	13%			
residual (including lighting)	34%	23%	33%	28%	47%	42%	18%	23%	23%	30%

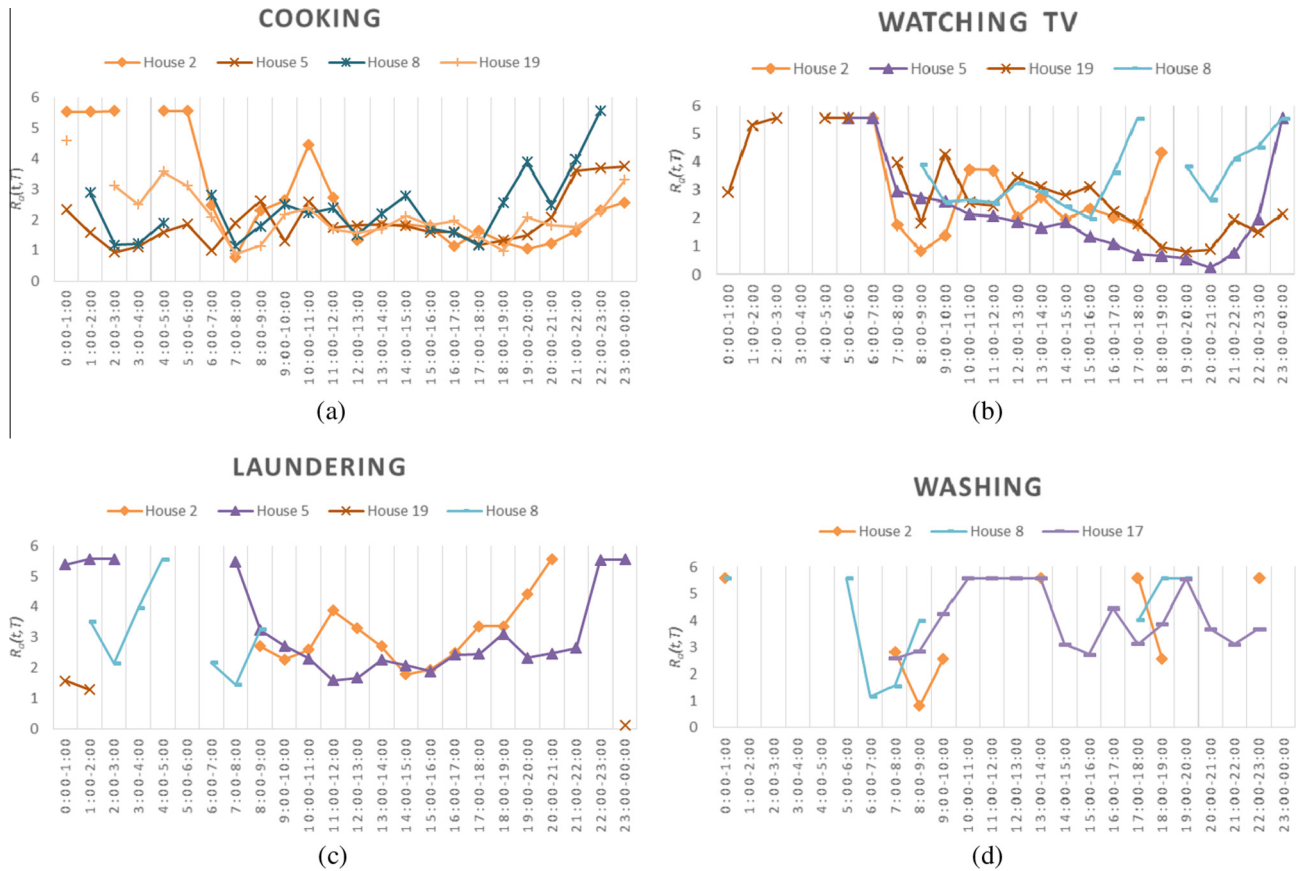


Fig. 6. Routine measure $R_a(t, T)$ for cooking, watching TV, laundering and washing activities for up to four households during each hourly slot t of a day (x -axis), for $T = 31$ days.

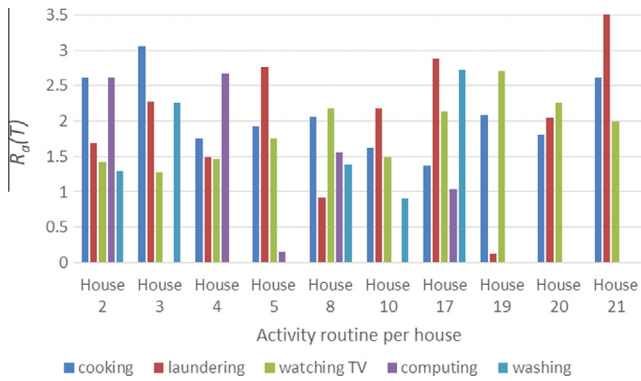


Fig. 7. $R_a(T)$ of five inferable activities. $R_a(T)$ denotes the 24-h average of $R_a(t, T)$, for $T = 31$ days computed using Eq. (3). Computing data was only available in a subset of households.

households have the smallest $R_{laundrying}(T)$ compared to other households. Both households are on the Economy7 tariff which has a lower off-peak cost between 10 PM and 8 AM [35]. Houses 2 and 8 are very consistent in their washing routine in the morning, unlike House 17, which has large variability throughout the day.

The average value of $R_a(T)$ over the course of a day, as shown in Fig. 7 for five activities for all 10 households, provides a rough measure of the overall routine in frequency and consistency of an activity. Fig. 7 enables us to understand and compare the degree of routine for each activity within and across households. For example, in Houses 2, 4 and 20 laundering and watching TV have similar

routines. Within House 19, only laundering is an activity that occurs at the same time, whereas cooking and laundering tend to vary more temporally.

5.4. Activities in households with similar composition

Household composition is an important and observable influence on the occurrence and time profile of domestic activities. We compare the energy intensity and routine of electricity-using activities within households of similar compositions. Specifically, we distinguish three types of household within our sample: families of four with two small children (2 households); families of four with two teenage children (3); pensioner couples (3).

From Table 8 we observe that there are no systematic differences in $EI_a(T)$ and $R_a(T)$ between household types. Even if energy intensities of the activities are similar, the routines for these activities can be very different (e.g., families with two teenage children). Despite similar occupancy, House 19 has a very different $EI_{laundrying}(T)$ and $R_a(T)$ to Houses 5 and 21. This could be due to their Economy 7 tariff. House 21 has the largest $R_{laundrying}(T)$, explained by the large variability in consistency and frequency of laundering. The $EI_{cooking}(T)$ and $R_{cooking}(T)$ for Houses 8 and 19 are similar despite the households having two and four members, respectively. However, there is a marked difference in their cooking patterns and demand profile. Fig. 8 shows that in the family of four with children (House 19) there is one clear peak for dinner in the evening, whereas the pensioner couple (House 8) spread out their cooking activity with multiple peaks throughout the day. The $R_{cooking}(T)$ value indicates that both households are consistent and frequent in their respective cooking routines. The plot also explains

Table 8

Energy intensity $EI_a(T)$ and routine $R_a(T)$ of three activities, namely cooking, laundering and watching TV, across households of similar composition, for $T = 31$ days.

	House	Total monthly electricity use (kWh)	$EI_{\text{cooking}}(T)$	$EI_{\text{laundering}}(T)$	$EI_{\text{watchingTV}}(T)$	$R_{\text{cooking}}(T)$	$R_{\text{laundering}}(T)$	$R_{\text{watchingTV}}(T)$
families of four with two small children	2	338	21%	4%	1%	2.61	1.68	1.43
	10	417	17%	6%	3%	1.62	2.18	1.49
families of four with two teenage children	5	636	15%	9%	2%	1.93	2.76	1.75
	19	248	15%	1%	2%	2.09	1.04	2.70
	21	333	23%	9%	1%	2.62	3.65	1.99
a pensioner couple	3	450	20%	12%	7%	3.05	2.27	1.28
	4	254	6%	3%	2%	1.75	1.49	1.47
	8	422	16%	4%	1%	2.06	0.92	2.18

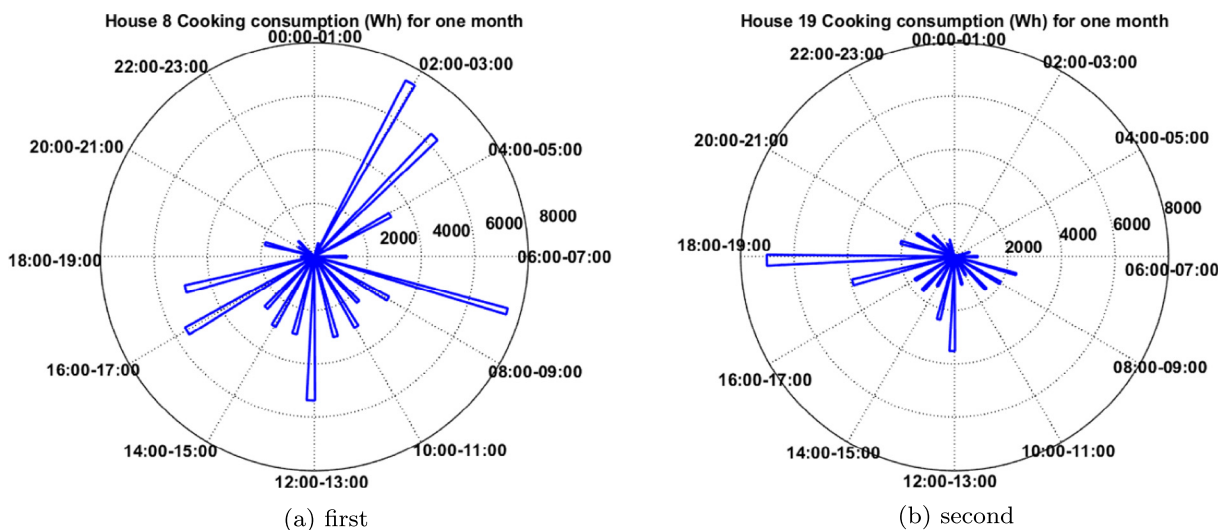


Fig. 8. Rose plots showing when cooking occurs and how much electricity it consumes over a month within each hourly slot (Houses 8 and 19). Total consumption is 68 kW h in House 8 and 37 kW h in House 19.

why House 8 consumes almost double the electricity for cooking compared to House 19.

6. Conclusions and future work

In this paper, we develop and demonstrate a novel methodology for inferring the energy and time use profile of domestic activities based on smart meter data and qualitative data obtained from interview and ethnography data. We implement the methodology on a data sensing and processing platform to analyse a month of data collected from ten households. We use the results to analyse the energy consequences of activities, and how they vary through time and between households. We provide an in-depth analysis of what can be inferred, given the qualitative and quantitative data available. We rely on real-time quantitative data from smart meters that collect active power measurements of individual appliances and/or aggregate active power for the whole household from which appliance usage information is extracted via non-intrusive load disaggregation approaches. These identify not only if an appliance was used, but also when the appliance started and finished its operation, as well as the electricity consumed for each run. We also rely on

qualitative data obtained from interview and ethnography data, which need only be collected at the beginning of the study and updated from time to time to reflect changes in household practices, to build an ontology mapping appliance end-use to activities.

Analysis of inferred activities shows that all energy intensive activities such as cooking and laundering can accurately be identified and that a significant portion of a household's electricity load can be attributed to the inferred activities. From a total set of 10 activities, a subset of 4–6 activities can be inferred for any given household accounting for 13–40% of total monthly demand. The remainder is accounted for by other electricity consumption by cold appliances, base load, and the residual including lighting. We define a metric to quantify the energy intensity of activities as a function of a household's total electricity demand, and a metric to indicate how routine an activity is through a 24-h daily cycle. We use these metrics to quantify the most energy intensive activities, and the consistency of their occurrence in time. This is useful for our primary intended application, which is household energy feedback.

Our work has several important applications. First, activities such as cooking, watching TV or hobbies, are meaningful ways in which householders spend their time at home. As such, providing

energy-related feedback to households through the lens of activities should help increase the salience of electricity consumption and its cost. Our activity-inference methodology enables this novel approach to energy feedback. Activity-itemized energy billing may present a viable alternative to more traditional aggregated billing practices. Second, linking energy-intensive activities to their time profile helps identify potentials for load shifting, demand side management, and identification of suitable tariffs. For example, understanding which households have less routine in activities may suggest the possibility of shifting these activities off-peak to minimise demand on generation and supply.

There can often be a complex mapping of appliance end-use onto activities in the ontology as well as uncertainties of inferring some activities, e.g., if an appliance end-use does not map directly onto the activity (e.g., socialising) or if an appliance end-use maps onto many potential activities (e.g., ICTs and the different computing-related activities). The importance of these uncertainties for household feedback is an important question, and something we plan to investigate further in subsequent research.

Our use of electricity data from smart meters means our methodology is potentially scalable alongside ongoing national rollouts of smart meters in the UK and elsewhere, supplemented by a few easy-to-use and low cost appliance-level monitors. However, our methodology as presented involves qualitative data as well as physical appliance surveys, which are used to build the activity ontologies per household. Future work will investigate a reduced-form methodology in which household interviews and video ethnography to identify appliance usage and activity routines are substituted by activity-based questionnaires that could be administered by remote or as part of a smart meter installation. In addition, appliance surveys which could be self-completed by households or carried out by smart meter installers with the households consent.

Acknowledgements

This work has been carried out as part of the REFIT project (Personalised Retrofit Decision Support Tools for UK Homes using Smart Home Technology', Grant Reference EP/K002368/1 and EP/K002430/1). REFIT is a consortium of three universities – Loughborough, Strathclyde and East Anglia – and ten industry stakeholders funded by the Engineering and Physical Sciences Research Council (EPSRC) under the Transforming Energy Demand in Buildings through Digital Innovation (BuildTEDDI) funding programme. For more information see: <http://www.epsrc.ac.uk> and <http://www.refitsmarthomes.org>. The data used in this paper can be obtained from <http://dx.doi.org/10.15129/9ab14b0e-19ac-4279-938f-27f643078cec>.

References

- Covrig CF, Ardelean M, Vasiljevskaja J, Mengolini A, Fulli G, Amoiralis E. Smart grid projects outlook 2014; 2014. <http://dx.doi.org/10.2790/22075>.
- Stephen B, Mutanen AJ, Galloway S, Burt G, Järventausta P. Enhanced load profiling for residential network customers. *IEEE Trans Power Delivery* 2014;29(1):88–96. <http://dx.doi.org/10.1109/TPWRD.2013.2287032>.
- Christensen TH, Friis F, Gram-Hanssen K, Plata A, Thronsen W. The role of households in the smart grid: a comparative study. In: ECEEE-2013 European council for an energy efficient economy 2013 summer study on energy efficiency. p. 2285–96.
- Erickson T, Li M, Kim Y, Deshpande A. The Dubuque electricity portal: evaluation of a city-scale residential electricity consumption feedback system. *Comput Syst* 2013;1203–12.
- Zeifman M, Roth K. Nonintrusive appliance load monitoring: review and outlook. *IEEE Trans Consum Electron* 2011;57:76–84.
- Hargreaves T, Nye M, Burgess J. Keeping energy visible? How householders interact with feedback from smart energy monitors in the longer term. *Energy Policy* 2013;52:1870–91.
- Wilson C, Hargreaves T, Hauxwell-Baldwin R. Smart homes and their users: a systematic analysis and key challenges. *Pers Ubiquit Comput* 2015;19:463–76.
- Gram-Hanssen K. New needs for better understanding of households energy consumption behaviour, lifestyle or practices? *Archit Eng Des Manage* 2014 (10):91–107.
- Katzeff C, Wangel J. Social practices, households, and design in the smart grid. In: *ICT innovations for sustainability*. Springer International Publishing; 2015. p. 351–65.
- Wilson C, Stankovic L, Stankovic V, Liao J, Coleman M, Hauxwell-Baldwin R, et al. Identifying the time profile of everyday activities in the home using smart meter data. In: *ECEEE-2015 European council for an energy efficient economy 2015 summer study on energy efficiency*.
- Office of National Statistics. National survey of time use. Tech rep. UK Government; 2000.
- Schwartz T, Stevens G, Jakobi T, Deneff S, Ramirez L, Wufli V, et al. What people do with consumption feedback: a long-term living lab study of a home energy management system. *Interact Comput* 2014.
- Widen J, Wckelgrd E. A high-resolution stochastic model of domestic activity patterns and electricity demand. *Appl Energy* 2009;87:1880–92.
- Durand-Daubin M. Household activities through various lenses: crossing surveys, diaries and electric consumption. In: *Behavior, energy and climate change conference*.
- Dunbabin P, Palmer J, Terry N. Electricity use by domestic appliances in English households. In: *ECEEE-2015 European council for an energy efficient economy 2015 summer study on energy efficiency*. p. 1439–50.
- Zoha A, Gluhak A, Imran MA, Rajasegarar S. Non-intrusive load monitoring approaches for disaggregated energy sensing: a survey. *Sensors* 2012;12:6838–66.
- Department of Energy and Climate Change. Smart metering equipment technical specifications version 2. Tech rep. UK Government; 2013.
- Zhao B, Stankovic L, Stankovic V. On a training-less solution for non-intrusive appliance load monitoring using graph signal processing. *IEEE Access* 2016;4:1784–99.
- Grunewald P, Layberry R. Measuring the relationship between time-use and electricity consumption. In: *ECEEE-2015 European council for an energy efficient economy 2015 summer study on energy efficiency*. p. 1439–50.
- Ellegard K, Jenny P. Visualizing energy consumption activities as a tool for making everyday life more sustainable. *Appl Energy* 2011(88):1920–6.
- Tanimoto J, Hagishima A, Sagara H. Validation of probabilistic methodology for generating actual inhabitants behaviour schedules for accurate prediction of maximum energy requirements. *Energy Build* 2008;40:316–22.
- Nguyenand TA, Aiello M. Energy intelligent buildings based on user activity: a survey. *Energy Build* 2013;56:244–57.
- Belley C, Gaboury S, Bouchard B, Bouzouane A. An efficient and inexpensive method for activity recognition within a smart home based on load signatures of appliances. *Pervasive Mob Comput* 2014(12):58–78.
- Clement J, Ploennigs J, Kabitzsch K. Detecting activities of daily living with smart meters. In: *Wichert R, Klausning H, editors. Ambient assisted living*. Berlin, Heidelberg: Springer-Verlag; 2014.
- Chen C, Barnan D, Cook D. Energy prediction based on resident's activity. In: *Int workshop knowledge discovery from sensor data*.
- Chen C, Cook D, Crandall A. The user side of sustainability: modelling behaviour and energy usage in the home. *Pervasive Mob Comput* 2013 (9):161–76.
- Cho H, Yamazaki T, Hahn M. AERO: extraction of users activities from electric power consumption data. *IEEE Trans Consum Electron* 2010(56):2011–8.
- Bates O, Clear A, Friday A, Hazas M, Morley J. Accounting for energy-reliant services within everyday life at home. *Pervasive Comput* 2012(7319):107–24.
- Froehlich J, Larson E, Gupta S, Cohn G, Reynolds M, Patel S. Disaggregated end-use energy sensing for the smart grid. *IEEE Pervasive Comput* 2011;1 (10):28–39.
- Clear A, Hazas M, Morley J, Friday A, Bates O. Domestic food and sustainable design: a study of university student cooking and its impacts. In: *ACM SIGCHI conference on human factors in computing systems*. p. 2447–56.
- Liao J, Stankovic L, Stankovic V. Detecting household activity patterns from smart meter data. In: *IEEE int conf intelligent environments*. p. 71–8.
- Stankovic L, Wilson C, Liao J, Stankovic V, Hauxwell-Baldwin R, Coleman M. Understanding domestic appliance use through their linkages to common activities. In: *EEDAL-2015, 8th int conf energy efficiency in domestic appliances and lighting*.
- Office of National Statistics & Eurostat. Survey on time use: activity coding list. Tech rep. UK Government; 2000.
- REFIT: Electrical Load Measurements (Cleaned) Dataset <<http://dx.doi.org/10.15129/9ab14b0e-19ac-4279-938f-27f643078cec>>.
- Murray D, Liao J, Stankovic L, Stankovic V, Wilson C, Coleman M, et al. A data management platform for personalised real-time energy feedback. In: *EEDAL*. p. 1–15.
- Liao J, Elafoudi G, Stankovic L, Stankovic V. Non-intrusive appliance load monitoring using low-resolution smart meter data. In: *Smart grid communications (SmartGridComm), 2014 IEEE international conference on*. p. 535–40. <http://dx.doi.org/10.1109/SmartGridComm.2014.7007702>.
- Hargreaves T, Wilson C, Hauxwell-Baldwin R. Control in the smart home. In: *DEMAND centre conference*.
- Murray D, Liao J, Stankovic L, Stankovic V. Understanding usage patterns of electric kettle and energy saving potential. *Appl Energy* 2016;171:231–42.
- Bertoldi P, Atanasiu B. Residential lighting consumption and saving potential in the enlarged EU. In: *4th international conference on energy efficiency in domestic appliances and lighting (EEDAL)*.