

Detecting Anomalies in Activities of Daily Living of Elderly Residents via Energy Disaggregation and Cox Processes

José Alcalá¹, Oliver Parson², Alex Rogers²

¹University of Alcalá, Madrid, Spain jmanuel.alcala@depeca.uah.es

²University of Southampton, Southampton, UK {osp,acr}@ecs.soton.ac.uk

ABSTRACT

Monitoring the health of the elderly living independently in their own homes is a key issue in building sustainable healthcare models which support a country's ageing population. Existing approaches have typically proposed remotely monitoring the behaviour of a household's occupants through the use of additional sensors. However the costs and privacy concerns of such sensors have significantly limited their potential for widespread adoption. In contrast, in this paper we propose an approach which detects Activities of Daily Living, which we use as a proxy for the health of the household residents. Our approach detects appliance usage from existing smart meter data, from which the unique daily routines of the household occupants are learned automatically via a log Gaussian Cox process. We evaluate our approach using two real-world data sets, and show it is able to detect over 80% of kettle uses while generating less than 10% false positives. Furthermore, our approach allows earlier interventions in households with a consistent routine and fewer false alarms in the remaining households, relative to a fixed-time intervention benchmark.

1. INTRODUCTION

Many countries facing a rapidly ageing population are grappling with the challenge of overhauling existing unsustainable healthcare models that seek to address illness but do little to promote wellness. One key area which has received significant attention is the aim of supporting and monitoring the health of the elderly living independently in their own homes.

To this end, many applications in the area of pervasive sensing have addressed the health monitoring of elderly people. Such approaches typically seek to remotely monitor an elderly person who is living alone, allowing rapid intervention if they deviate from their normal routines [Scanaiil *et al.*, 2006]. For example, a lightweight system might use a personal alarm worn by the elderly resident with a button that is pressed in an emergency, while a more intrusive system might require the use of closed circuit television or webcams to monitor the elderly resident. However, these systems typically require either the installation of additional costly sensors or a change to the elderly resident's routine. Such invasive moni-

toring also raises concerns regarding the privacy and dignity of the elderly, and as such, proposed systems typically seek to use indirect forms of monitoring.

Now, in many countries, the widespread installation of smart electricity meters is well underway. Such meters provide measurements of household electricity consumption, across the network, at temporal resolutions ranging from hourly measurements to sub-minute level monitoring. Furthermore, recent work within the artificial intelligence literature has made significant process developing algorithms to disaggregate a single measurement of electricity consumption into individual appliance usage. As such, the use of smart meter data to inform a pervasive health monitoring system is an attractive solution, since it requires no additional sensor deployment, and could be provided at little additional cost by the energy retailer.

To address these shortcomings, in this paper we present an approach which uses aggregate electricity data provided by existing and planned deployments of smart meters to monitor the Activities of Daily Living of elderly residents in their own homes, which we use as a proxy to the elderly resident's health. We use standard disaggregation approaches to recognise and detect commonly used appliances. Our approach models daily appliance usage using a periodic log Gaussian Cox process, allowing a usage schedule to be learned automatically for each individual household. This model is then used to predict appliance usage along with a measure of uncertainty, allowing unusual usage patterns to be detected, and easily interpreted alerts to be generated which can be used to precipitate an intervention (e.g. warning that on 90% of previous days the kettle would have been used by this time).

The focus of this work is to produce an approach which can be applied to any UK household with a smart meter installed. Crucially, our approach requires only 10 seconds aggregate power data as is available via the in-home ZigBee network from UK smart meters [Department of Energy & Climate Change, 2013], and does not require information regarding which other appliances are present in the household, nor does it require a distribution of the appliance's usage to be specified a-priori. We focus on the kettle in the evaluations presented in this paper, since the appliance is present in almost all UK households and also forms a key part of the daily routine of the elderly residents in the UK.

Our contributions are summarised as follows:

- We describe a disaggregation method capable of detecting uses of specific appliances from aggregate smart meter data. We employ the difference hidden Markov model (dHMM)

to disaggregate appliances from the aggregate load, in which the differences between consecutive aggregate power measurements are used as the observation sequence in the hidden Markov model (HMM). However, the dHMM is allowed to filter observations with low likelihood given the appliance model, and as such only the step changes matching the appliance of interest are used to estimate the appliance's states. In addition, the total power demand is also used as a secondary observation sequence in the HMM to ensure the appliance is never estimated to draw more power than the household aggregate power demand.

- We present a novel method for detecting deviations from the standard routine of a household. As such, the unique daily and weekly routine for any household can be learned given a history of smart meter data. Occurrences of individual appliances over the day and week are modelled as a Poisson distribution where the intensity varies over time. This intensity function is learned using a log Gaussian Cox process, where the log of the intensity function is modelled using a Gaussian process. By using two periodic-exponential kernels, our approach ensures that the likelihood of usage varies smoothly throughout the day and week, while also wrapping continuously across the midnight and end of week boundary. This household-specific usage model can then be used to calculate the cumulative likelihood that the modelled appliance would have been used by any point in a day to a given level of certainty.
- We provide a thorough evaluation of both the disaggregation and usage modelling components of our approach. The disaggregation approach is evaluated on the task of detecting kettle usage, using a UK-based data set, which crucially contains high-resolution 1 second aggregate electricity data (which we downsample to 10 second data to mimic the reporting rate of UK smart meters) as well as ground truth electricity consumption of the kettle. We show that over 80% of kettle uses can be correctly detected while also detecting less than 10% false positives. The usage model is then evaluated using the Household Electricity Survey data set,¹ which contains low-resolution kettle usage data from a number of elderly single occupant households. We show that our approach allows earlier interventions in households with a consistent routine (25 out of 32 households) and fewer false alarms in the remaining households, relative to a fixed-time intervention benchmark.

In the following sections, we first discuss recent work related to remote health monitoring and energy disaggregation (Section 2). We then describe the appliance detection and routine modelling components of our approach (Section 3), followed by a thorough evaluation of each component using two different data sets (Section 4). Finally, we conclude and give directions for future work in Section 5.

2. RELATED WORK

In this section, we first discuss traditional methods for health monitoring through pervasive sensors, before providing an overview of the state of the art in energy disaggregation, and its application to the domain of health monitoring.

¹<https://www.gov.uk/government/collections/household-electricity-survey>

2.1 Health Monitoring

Traditional methods for monitoring the health of elderly adults typically involve either daily visits by a care worker to the elderly person's house or requiring the elderly person to move to a care home. However, the former requires a significant amount of resources for each household, and as a result each household can only be visited for a short amount of time each day. The latter allows the health of the elderly person to be monitored constantly, but such care homes are costly and require the elderly person to be uprooted from their own home. Neither approach allows the elderly resident to maintain their independence, dignity and privacy while also receiving constant monitoring, and thus motivates the use of smart home technologies to provide low-cost and non-intrusive approaches to monitoring the health of elderly residents living in their own homes. Such smart home technologies can be classified as either direct or indirect forms of health monitoring, which we discuss further in this section.

Direct health monitoring refers to the measurement of health parameters (e.g. blood pressure) by integrating such sensors in to normal household structures. For example, Togawa [1998] suggest that blood pressure can be measured automatically via sensors in a toilet seat, or blood samples can be taken by a device mimicking a mosquito. Although such solutions have the potential to accurately identify changes in an individual's health parameters, they are clearly expensive and intrusive and as a result do not lend themselves well to affordable and scalable deployments.

Indirect health monitoring aims to infer the health of the household occupants through the use of pervasive sensors. Such approaches typically aim to determine whether Activities of Daily Living (ADLs) are being carried out as part of normal routines, such as feeding, bathing, dressing etc. However, while many ADLs will be common to many households, the exact routine is likely to be unique to a single household. Lim *et al.* [2008] propose a daily activity recognition system for elderly people which uses floor pressure sensors to detect deviations from normal routine behaviour. Ohta *et al.* [2002] propose a system in which infrared sensors are deployed in every room, from which the health condition of the household occupants is estimated by comparing the duration of stays in specific rooms. Ross [2004] discuss the use of Motes, a sensor developed by Intel, designed to be attached to everyday household items, such as clothes, cups or medicine containers. Such sensors collected data from which ADLs can be inferred. However, all such approaches require the installation of additional sensors within the home, which severely limits the scalability of such solutions. This motivates the use of data from existing sensors, such as smart meters, which we discuss in the following section.

2.2 Energy Disaggregation

Energy disaggregation is the process of estimating the usage of individual appliances in a home from a single point of measurement [Hart, 1992], such as a smart meter. Hidden Markov models (HMMs) have become a popular tool for modelling appliances in recent years [Zia *et al.*, 2011; Kolter and Johnson, 2011; Kim *et al.*, 2011; Kolter and Jaakkola, 2012; Johnson and Willsky, 2013; Zoha *et al.*, 2013; Zhong *et al.*, 2014], which are typically combined to form a factorial HMM representing most or all appliances in the household. Such factorial HMMs are then used to disaggregate the total electricity consumption of a household into individual appliances. Since such models require that the sum of the power of the modelled appliances approximately sum to the mea-

sured aggregate demand, it is necessary to model many appliances in each household. However, it is often impractical to model such a large number of appliances when sub-metered data is not available from each individual household, and moreover it is not even known which appliances are present in each household. As such, more practical approaches have aimed to develop one-vs-all disaggregation methods, which crucially only require the appliance of interest to be modelled.

To the best of our knowledge, the only use of the disaggregation of smart meter data for health monitoring was carried out by Song *et al.* [2014]. The authors use a one-vs-all iterative time-dependent hidden Markov model to disaggregate five appliances from the aggregate load. However, the approach was only evaluated using a single household, so the generality of the approach is not known. Furthermore, their approach relies on appliance time of use distributions, for which the detailed operational parameters of each appliance studied were specified a-priori, rather than learned from data. As a result, the approach could not be deployed at scale without a detailed survey to be carried out in each household first; significantly limiting its potential for widespread adoption.

Usage pattern modelling is also a key factor when detecting deviations from normal daily routines. In the aforementioned study, Song *et al.* [2014] use the ADL profile to determine a usage pattern model for each appliance. This usage pattern is essentially a histogram with hourly information about the frequency of use for each appliance. However, the usage patterns are assumed to be the same for every day of the week. Similarly, Zhang *et al.* [2014] obtain the daily usage pattern by grouping sequences of on-off events. For each appliance and each day, the authors pair on and off events and extract three features for the usage pattern: timestamp of an event, duration of the paired event and number of paired events. The events are then clustered and classified according to a heuristic threshold. Again, this approach ignores day of week information. A different approach is adopted by Chen *et al.* [2013], in which usage patterns do not depend on the time of the day but on the Activity of Daily Living the tenant is performing. The authors apply Latent Dirichlet Allocation to associate ADLs with usage patterns and, for each ADL, give a probability and time of use of each appliance. This study does not show results for detecting deviations from standard patterns so it is difficult to evaluate its performance. Besides, the approach requires the time-consuming task of collecting ADLs in each house and their experiments are carried out over families instead of single elderly households.

3. MODEL

In this section, we first describe our disaggregation algorithm which detects uses of appliances from 10 second aggregate power data available from smart meters (Section 3.1). Although not the core novelty of this work, the disaggregation is a necessary component to extract appliance usage data, which we then model using a log Gaussian Cox process to learn the daily and weekly patterns of appliance usage for individual households (Section 3.2).

We use the kettle as an example appliance throughout this section and our evaluation in Section 4. The reasons for this choice are threefold. First, kettles are present in almost all UK households, as shown by Figure 1 which shows the 30 most common appliances in UK homes, taken from the Household Electricity Survey [Zimmermann *et al.*, 2012]. Second, kettles are used regularly and form a core part of the daily routine of UK households, and especially that of the elderly residents as discussed in Section 3.2. Third, kettles




Figure 1: Appliance ownership in the UK.

exhibit regular power consumption since the manufacturers typically seek to set the power of the appliance at the upper end of the feasible range (as typically defined by the maximum fuse rating allowable under building regulations). It is worth noting that, although our approach to disaggregation and usage modelling could be applied to other appliances, none provide the same combination of benefits of the kettle. For example, fridges and freezers are present in many households and are relatively easy to detect [Parson *et al.*, 2014], although their automatic operation mean their usage is not indicative of the routines of the household occupants. Alternatively, washing machines are also present in many households and easy to detect, though they are used relatively infrequently and thus are less useful when providing early interventions when normal routines are not followed.

3.1 Appliance Detection for Smart Meter Data

We adopt the disaggregation model presented by Parson *et al.* [2012], which consists of a difference hidden Markov model in which the latent variables correspond to the appliance states and the observed variables correspond to the difference between consecutive readings of the aggregate power demand. In addition, unlike a normal HMM, only the observations which can be plausibly explained by the model are retained, allowing step changes corresponding to appliances other than the one of interest to be filtered out. Furthermore, in order to prevent the model from estimating that a high power appliance has remained on in the case that the off event has been missed, the aggregate power demand is used as an additional constraint on the appliances that could be on at any given point.

Figure 2 shows a Bayesian network representation of the model. In the figure, each latent discrete variable, z_t , in the Markov chain represents the state of the appliance at an instant in time. Each variable z_t takes on an integer value in the range $[1, K]$ where K is the number of states.

In a standard HMM, each variable in the Markov chain emits a




Figure 2: A difference HMM variant where observation y_3 shown by dashed lines has been filtered out.

single observation. However, in our model we consider two observation sequences \mathbf{x} and \mathbf{y} . Sequence \mathbf{x} corresponds to the household aggregate power demand measured by the smart meter. These aggregate power observations are used to restrict the time slices in which an appliance can be on to only those when the aggregate power demand is greater than that of the individual appliance. Sequence \mathbf{y} is derived from \mathbf{x} , and corresponds to the difference between two consecutive aggregate power readings such that $y_t = x_t - x_{t-1}$ (hence this model is referred to as a difference HMM). These derived observations are used to infer the probability that a change in aggregate power, y_t , was generated by two consecutive appliance states.

The dependencies between the variables in our graphical model can be defined by a set of three appliance parameters: $\theta = \{\pi, \mathbf{A}, \phi\}$, respectively corresponding to the probability of the appliance's initial state, the transition probabilities between states and the probability that an observation was generated by an appliance state. We now define the functions governed by these parameters.

The probability of the appliance's starting state at $t = 1$ is represented by the vector π such that:

$$p(z_1 = k) = \pi_k \quad (1)$$

The transition probabilities from state i at $t - 1$ to state j at t are represented by the matrix \mathbf{A} such that:

$$p(z_t = j | z_{t-1} = i) = A_{i,j} \quad (2)$$

We assume that each appliance has a Gaussian distributed power demand:

$$w_t | z_t, \phi \sim \mathcal{N}(\mu_{z_t}, \sigma_{z_t}^2) \quad (3)$$

where w_t is the appliance's power demand.

The emission probabilities for \mathbf{x} are described by a function governed by parameters ϕ , which in our case are assumed to be Gaussian distributed such that:

$$y_t | z_t, z_{t-1}, \phi \sim \mathcal{N}(\mu_{z_t} - \mu_{z_{t-1}}, \sigma_{z_t}^2 + \sigma_{z_{t-1}}^2) \quad (4)$$

where $\phi_k = \{\mu_k, \sigma_k^2\}$, and μ_k and σ_k^2 are the mean and variance of the Gaussian distribution describing this appliance's power draw in state k . Equation 4 is used to evaluate the probability that a change in the aggregate power was generated by an appliance transition between two states.

Equations 1, 2 and 4 are the minimum definitions needed to define a difference HMM. However, by using the change in aggregate power as the only observation sequence, the model does not impose the constraint that appliances can only be 'on' when the observed aggregate power is greater than the appliance's power. We impose this constraint by considering the aggregate power demand, x_t , as a censored reading of an appliance's power demand, w_t , which we incorporate into our model using an additional emission function representing the cumulative distribution function of an appliance's Gaussian distributed power demand:

$$\begin{aligned} P(w_{z_t} \leq x_t | z_t, \phi) &= \int_{-\infty}^{x_t} \mathcal{N}(\mu_{z_t}, \sigma_{z_t}^2) dw \\ &= \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{x_t - \mu_{z_t}}{\sigma_{z_t} \sqrt{2}} \right) \right] \end{aligned} \quad (5)$$

where $\operatorname{erf}()$ represents the Error Function [Andrews, 1985].


This emission function constrains the model such that when the aggregate power reading is much less than the mean power draw of an appliance state, then the probability of that state will tend towards 0. However, if the aggregate power reading is much greater than the mean power draw of the appliance state, the emission probability of that state will tend towards 1. This ensures that the appliance of interest is never estimated to be on if its power demand is much larger than the households aggregate power demand.

3.2 Usage Pattern Monitoring via Log Gaussian Cox Process

Our approach exploits the routines which are followed by elderly single-occupant UK households. However, it is important to note that although such routines are relatively constant in the long-term, they are also unique to each individual household, and as such must be learned individually. Figures 3 and 4 show histograms of kettle usage across a day, week and year for two households. It should be noted that the gap in the histogram in Figure 3 (c) is due to missing sensor data. It can be seen that both households exhibit a clearly daily routine, though both routines are subtly different. In addition, it can be seen that the house in Figure 3 shows no difference in routine between different days of the week, while the house in Figure 4 shows a clear difference, with frequent kettle usage from Monday to Thursday, and reduced usage from Friday to Sunday. From a manual inspection of the data, we found that 31 out of the 32 houses present some daily periodicity, while 25 out of 32 houses present weekly periodicity. Last, similar to the daily routines, it can be seen that the kettle usage varies across the year for both households, although again the routine is unique for each household. It is exactly these routines which our approach learns for individual households using kettle uses detected from smart meter data.


We model the usage of individual appliances (and specifically, the kettle in this study) in a household using a log Gaussian Cox process. This model is chosen as it allows arbitrarily complex usage patterns to be learned from point data. Furthermore, it provides a principled model for encoding daily and weekly periodicity in the usage model. The model also provides estimates of the uncertainty around each prediction. It is important to bear in mind that the choice of the appliance is crucial to gain an insight into the human behaviour behind the usage pattern. The appliance should be one that is manually operated, and also exhibits daily and predictable use. We now describe the model in more detail.

We map the list of appliance uses for a given household, as pro-



(a) Frequency of usage per hour of day. (b) Frequency of usage per day of week. (c) Frequency of usage per week of year.

Figure 3: Frequency of kettle usage in house 101017 of HES data set.



(a) Frequency of usage per hour of day. (b) Frequency of usage per day of week. (c) Frequency of usage per week of year.

Figure 4: Frequency of kettle usage in house 102003 of HES data set.

duced by the disaggregation model, such that only time of week information is retained. This produces a temporal point pattern covering a 7 day period, as shown in Figure 5. We assume that the temporal distribution of these points is generated by a random process with an intensity function $\lambda(x)$, where x is a continuous variable representing a time throughout the week. Furthermore, if we bin the data into fixed length bins, the expected number of appliance uses is $E[\lambda(X)]$ and the number of points within that bin is distributed by a Poisson distribution [Cressie and Wikle, 2011].

$$\rho(E[\lambda(X)]) = \text{Poisson}\left(\int_X \lambda(X)\right) \quad (6)$$

In fact, the non-homogeneity of the process over the whole week is due to a time-dependency of λ that follows an unknown function and is itself generated by another stochastic process. This double stochastic process is known as a Cox process [Cox and Isham, 1980]. Therefore, it is necessary to learn the hyperparameters of this second stochastic process. To do so, we use a non-parametric statistical method referred as the log Gaussian Cox process. This method uses a Gaussian process to learn the temporal point processes [Møller *et al.*, 1998]. In this case, the log of the intensity function is assumed to be generated by a Gaussian process:

$$z(x) = \log(\lambda(x)) \sim \text{GP}(m(x), K(x, x')) \quad (7)$$

where:

$$m(x) = E[z(x)] \quad (8)$$

$$K(x, x') = \text{Cov}(z(x), z(x')) \quad (9)$$

Equations 8 and 9 are the mean and the kernel function of a Gaussian process and describe the Gaussian process regression, while x is a variable in the temporal points and x' refers to the prediction point. We also want to use a kernel that ensures the model predictions will change smoothly over the course of the day, while also allowing both the daily and weekly periodicity to be explicitly encoded in the model. To do so, we use two periodic-exponential kernels, each of which have three hyperparameters: (i) the variance, σ , that represents the strength of the dependency over other temporal points, (ii) the lengthscale, ℓ , that represents a window in which the points have the greatest effect on the regression and (iii) the period, ρ , which represents the daily or weekly repeating patterns in the temporal point process:

$$K(x, x') = \sigma^2 \exp\left(\frac{-2\sin^2(\pi|x - x'|/\rho)}{\ell^2}\right) \quad (10)$$

It should be noted that yearly periodicity (e.g. seasonal variations) could also be trivially included in this model. However, the HES data set (which we use for our evaluation in Section 4.2) only monitors three elderly single occupant households for one year, while the remaining households were monitored for one month. As such, we focus on only daily and weekly repeating patterns in this paper.

The next step is to obtain a posterior distribution $\rho(z|X_i)$ to allow predictions to be made. This is a computationally complex problem which requires the integration of the Gaussian prior (Equation 7) over a Poisson likelihood distribution. We solve this equation using the Laplace approximation [Friston *et al.*, 2007], which provides an efficient approximation at a small cost in accuracy. We implement the log Gaussian Cox process using the GPy Gaussian Process framework in python [The GPy authors, 2012 2014].

Figure 6 (a) shows the posterior distribution of kettle usage throughout the week, while Figure 6 (b) shows the covariance function of the relationship between midnight on Monday and each other point during the week. It can be seen that points close to each other in time of day share a strong correlation via the exponential part of the kernel, while points that occur at a similar time but on different days also share strong correlation via the first periodic kernel. Similarly, points that are separated by exactly one week also share a strong correlation via the second periodic kernel.

We initialise the kernel with a lengthscale of six hours and a variance of one, before optimising these parameters given the disaggregated appliance usage data, while holding the periodicity constant at both one day and one week. This ensures that both the model parameters and the regression fit is learned for each household. Figure 7 (a) shows an example regression after optimising the model parameters and returning the exponential shape to the intensity function for the kettle in a single household. The figure shows the mean of the regression as well as the 95% confidence interval. It can be seen that the usage pattern for this household shows a high chance of the kettle being used at around 7am when the occupant wakes up and another at roughly 9pm before the occupant goes to bed, and also three smaller peaks during the mid-morning, lunchtime and mid-afternoon periods.

Finally, we transform the regression fit to represent the cumulative probability that an appliance, a , was used at least once by time of day, t . This is calculated using:

$$p(a_{0:t}) = 1 - [(1 - p(a_{0:t-1})) (1 - p(a_t))] \quad (11)$$

Figure 7 (b) shows an example of this cumulative function for the same household as in Figure 7 (a) over the course of one day. It can be seen that while the probability of kettle usage within each interval is not necessarily high, the probability of a single usage quickly accumulates throughout the day. For example, the graph shows that there is a 80% chance of the kettle being used at least once by 11am, which could be a favourable threshold for intervention in the case that the kettle has not been used by this point of the day. Our approach applies this threshold to all households, in order to obtain a unique intervention time for each household. As a result, it is possible to recommend an intervention in any household which would have likely seen the modelled appliance being used by that time of the day, yet no usage was detected by the disaggregation algorithm.

4. EMPIRICAL EVALUATION

We now evaluate our model using the kettle as the appliance to be disaggregated and modelled via a log Gaussian Cox process for the same reasons discussed in Section 3. We first provide an evaluation of the disaggregation model using the Colden Common data set (Section 4.1), followed by an evaluation of the log Gaussian Cox process model using the HES data set (Section 4.2).

4.1 Appliance Detection for Smart Meter Data

We evaluate the disaggregation of kettle uses using a data set collected from the village of Colden Common in the UK. This data set contains both household aggregate power data (disaggregation input) and kettle power data (ground truth) for 13 houses. This data set was particularly well-suited to our evaluation as the aggregate power data was collected at 1 second intervals, which we down-sampled to 10 second intervals to match the reporting rate of UK smart meters over the in-home network [Department of Energy & Climate Change, 2013].




Figure 5: Occurrences and histogram of kettle usage during a week.

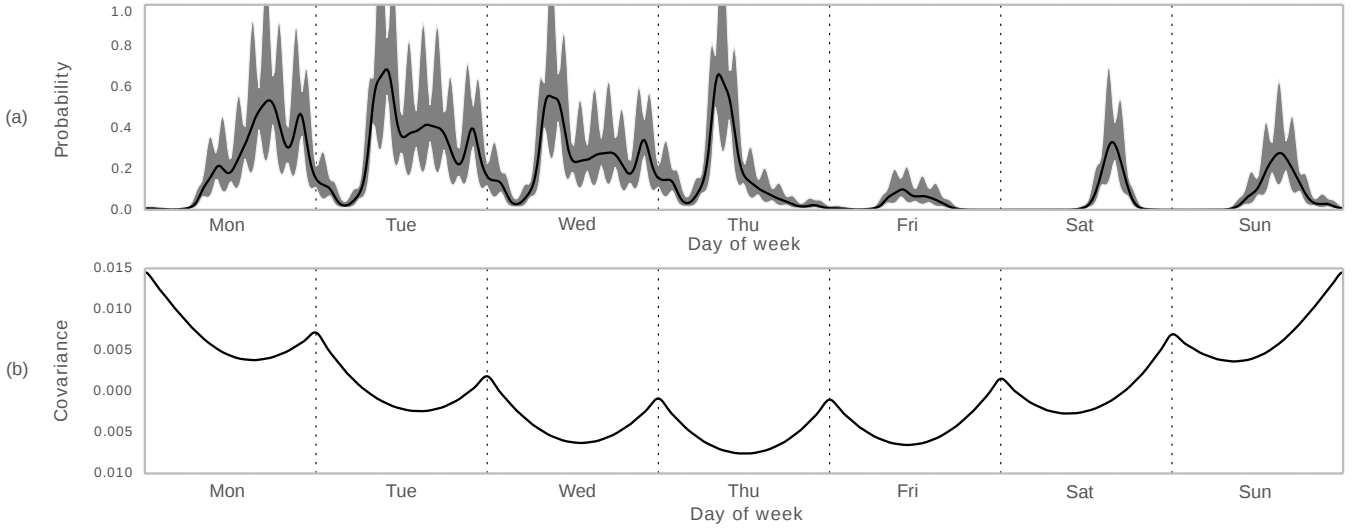


Figure 6: (a) The probability of a kettle being used at each point in the week for a single household. (b) The covariance function between each point in the week and midnight on Monday.

We measure the accuracy of kettle detection using the receiver operating characteristic (ROC). By plotting a ROC curve as some threshold of the model is varied, the trade-off between true positive rate (TPR) and false positive rate (FPR) can be investigated. In our deployment, the true positive rate refers to the fraction of actual kettle uses which were successfully detected, while the false positive rate refers to the fraction of times the kettle was incorrectly classified as being used out of all bins when the kettle was not actually used. On a ROC curve, a TPR of 1 and a FPR of 0 indicates perfect classification, while the $TPR = FPR$ diagonal represents random classification. It is necessary to bin the data such that the classifications correspond to individual uses of the kettle rather than instants in time during which the kettle was on or off. However, a small bin size (e.g. 1 minute) will result in a large amount of true negatives due to the relatively rare occurrence of the kettle being used, while a large bin size (e.g. 6 hours) could potentially result in incorrect detections of the kettle occurring in the same bin as an actual usage even if the detection was missed by a number of hours. For this reason, we use a bin size of 30 minutes to provide a reasonable trade-off between these two extremes.

We compare our approach against a naive Bayes classifier with one class, similar to that of Hart [1992]. In this approach, posi-

tive step changes in aggregate power greater than 50 W are paired with subsequent negative step changes when the absolute difference between the step changes is less than 100 W and the duration separating the step changes is less than 10 minutes. The positive and negative step changes, along with the duration separating the two step changes, are then used as features in a one-class naive Bayes classifier, in which the product of the independent probabilities of each feature is compared against a likelihood threshold. This likelihood threshold is varied to provide different trade-offs between TPR and FPR. We use a single general appliance model for the kettle for both the naive Bayes classifier and the dHMM methods due to the similarity between the electrical signature of UK kettles.

Figure 8 shows a ROC curve which compares the accuracy of the dHMM and the naive Bayes classifier. The curve is made up of a set of points, each of which represent the mean TPR and FPR over each household in the data set as the two model thresholds are varied. It can be seen that the dHMM provides highly accurate disaggregation, with the curve outperforming the performance of the naive Bayes classifier across all trade-offs between TPR and FPR. A particularly favourable trade-off between TPR and FPR can be seen, in which more than 80% of the kettle uses were correctly detected while also detecting less than 10% false positives.




Figure 7: (a) Learned probability that the kettle would be used in a half hour interval. (b) Cumulative probability the kettle would be used at least once throughout the day.

4.2 Usage Pattern Monitoring via Log Gaussian Cox Process

The Colden Common data set could not be used to evaluate the usage pattern model as it does not contain metadata identifying which households were occupied by an elderly resident living independently. Instead, we evaluate the monitoring of usage patterns using the Household Electricity Survey (HES) data set [Zimmermann *et al.*, 2012]. Similar to the Colden Common data set, the HES data set also contains kettle electricity data, but covers 250 households in the village of Colden Common, Hampshire, UK, of which 32 households are known to be occupied by elderly residents living alone. It is worth noting that the HES data set could not also be used for evaluating our disaggregation algorithm in Section 4.1 as it contains cumulative energy data over 2 minute intervals, rather than the 1 second instantaneous power data available in the Colden Common data set. For these reasons, it was necessary to evaluate the disaggregation and usage modelling separately, although it is important to note that in a real deployment, disaggregation errors would propagate through to the usage modelling and appear as a false positive or false negative of an intervention. We have focused our evaluation on single occupant elderly households for two reasons: (i) the usage pattern of a single occupant is more stable and predictable than several people living together, and (ii) it is crucial to intervene as early as possible when elderly people are living alone as no other monitoring is possible.

For each household, we extract the usage pattern as described in Section 3.2. The same kernel was used for all households, except for the hyperparameters (variance and lengthscale) which were optimised for individual households while the periodicity was constrained to one day and one week. We use a bin size of 30 minutes given the limited usage data that was available.

Figure 9 shows the false alarm rate against the hour of interven-

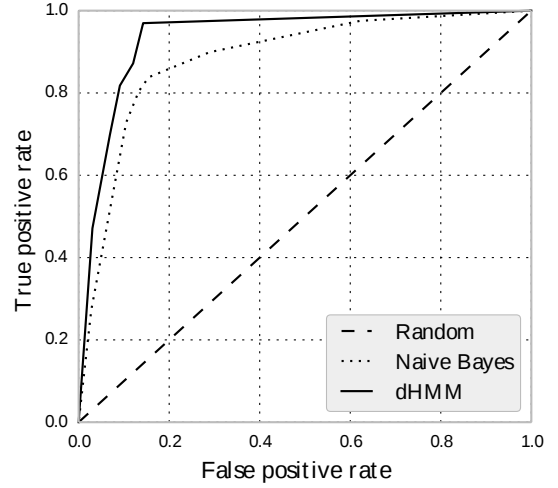


Figure 8: ROC curve showing kettle disaggregation accuracy.




Figure 9: False positive count for various intervention times. Each line represents one day of week for one house.

tion for our approach for each day of week for each household. We define a false alarm at a given time of day as a recommended intervention in a household which has not used the kettle by that time but uses the kettle at a later point during that day. Conversely, a correctly raised alarm is a recommended intervention for a given day in which the kettle is not used at all. As such, interventions earlier in the day allow the elderly resident to receive care as soon as possible but can only be achieved at the cost of a high false alarm rate, while interventions later in the day are necessary to achieve a low false alarm rate. As a result, the hour of day of the intervention can be determined for each individual house given an acceptable false alarm rate. It can be seen that the false alarm rate drops rapidly to below 0.1 (corresponding to a single false alarm every 10 days) for most houses and days of week between the hours of approximately 5am and 5pm. However, an early intervention within this 12 hour window could be crucial to ensure that an elderly person receives the attention they need as soon as possible. It is exactly this characteristic that our approach exploits, allowing early interventions in households when a low false alarm can be maintained, while allowing later interventions in houses which use the kettle less frequently or later in the day.

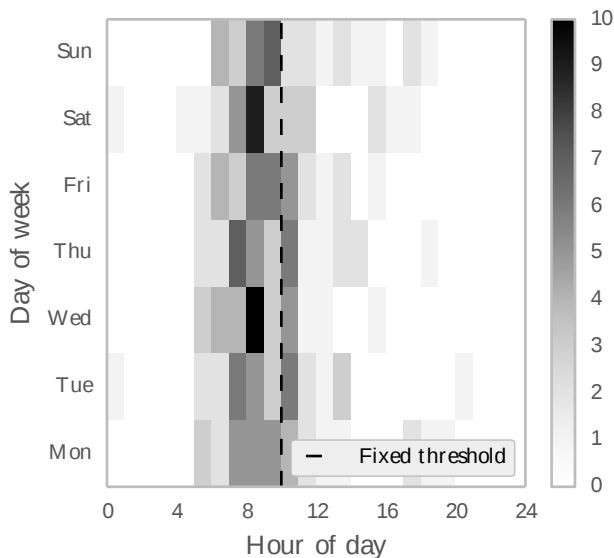


Figure 10: Intervention times of our approach (LGCP) against fixed-time intervention benchmark.

We compare our approach against a fixed threshold method of intervention, which will intervene in any household if there has not been any use of the disaggregated appliance before a fixed time of day. This benchmark was chosen as it represents the optimal intervention point given no individual information about household routines. We optimised the intervention time to achieve the earliest possible intervention while averaging a false alarm rate of 10% across all households.

Figure 10 shows the intervention time of our approach against the fixed-time benchmark across all single occupant elderly households in the HES data set. On average, our approach allows earlier interventions in 25 out of the 32 households, relative to the fixed-time benchmark. Furthermore, our approach intervenes later in the remaining households, in order to ensure the tolerable false alarm rate of 10% in every household, rather than on average across all households as is achieved by the benchmark. As a result, our approach is able to identify the households which really need an intervention, as opposed to those that are just following their normal schedule of using the kettle later in the day.

Finally, it is important to note that our approach aims to provide intuitive information to a friend, relative or health care worker. For example, the system could produce an alert such as: ‘on 90% of days, the kettle would have been used by now in this household.’ Such alerts could be delivered by text message, allowing the recipient to make an informed decision on what action to take. For instance, if the text message is received by a relative who had taken the elderly resident out for lunch, there is relatively little cost of generating this false positive. Alternatively, if the text message is received by a care worker, the care worker could make a phone call to the elderly occupant to check everything is okay, and raise the case as an emergency if no contact can be made. Such a system clearly requires less resources to monitor elderly residents living independently than making daily phone calls or in-person visits to all such households, as interventions only occur due to unusual deviations from standard routines.

5. CONCLUSION

In this paper we have presented a novel approach to monitoring the health of elderly occupants who live alone. Our aim is to provide basic non-intrusive health monitoring which can be deployed at a large scale, which crucially requires no additional sensors to be installed within each household. Our approach is able to identify appliance activities from smart meter data and extract the pattern of usage, which is used to monitor the health of the household’s occupant. Through our proposed approach, the independence of elderly people can be maintained while non-intrusively monitoring their activity.

In a detailed empirical evaluation, we applied our approach to the detection of the kettle in a 13 UK homes. Our algorithm detected over 80% of the times the kettle was used while generating less than 10% false positives, using a general model for the kettle which applies to all households without any further tuning. We then model the usage of the appliance using a log Gaussian Cox process, and show that our algorithm is able to learn the usage pattern unique to each household and efficiently trigger an alarm when any deviation from this pattern is detected by not using the specific appliance. Our evaluation indicates that our approach allows earlier interventions in households with a consistent routine (25 out of 32 households on average) and fewer false alarms in the remaining households, relative to a fixed-time intervention benchmark. Furthermore, the use of a log Gaussian Cox process allows the usage pattern to change over time as well as providing a measure of uncertainty.

There are a number of potential directions for future work. Having shown that the kettle can be both detected and used as a proxy for routine in UK households, we would like to investigate which other appliances can also be disaggregated and used to model routines in other countries. Furthermore, in this paper we have investigated a strategy which aims to intervene if a modelled appliance is not used at all, however the same usage pattern model could be applied to detect whether an appliance was being used but at unusual times of the day. In addition, we would like to explore incorporating strong priors into the structure of the log Gaussian Cox process to handle sparse initial training data. Finally, we are also working with our energy company partner to initiate a trial deployment of our approach to gain further insight into the needs and challenges in this important problem domain.

References

- Larry C Andrews. *Special functions for engineers and applied mathematicians*. Macmillan, 1985.
- Yi-Cheng Chen, Yu-Lun Ko, Wen-Chih Peng, and Wang-Chien Lee. Mining appliance usage patterns in smart home environment. In *Advances in Knowledge Discovery and Data Mining*, volume 7818 of *Lecture Notes in Computer Science*, pages 99–110. Springer Berlin Heidelberg, 2013.
- David Roxbee Cox and Valerie Isham. *Point processes*, volume 12. CRC Press, 1980.
- Noel Cressie and Christopher K Wikle. *Statistics for spatio-temporal data*. John Wiley & Sons, 2011.
- Department of Energy & Climate Change. Smart Metering Equipment Technical Specifications Version 2. Technical report, UK, 2013.

- Karl Friston, Jérémie Mattout, Nelson Trujillo-Barreto, John Ashburner, and Will Penny. Variational free energy and the laplace approximation. *Neuroimage*, 34(1):220–234, 2007.
- G. W. Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870–1891, 1992.
- Matthew J Johnson and Alan S Willsky. Bayesian Nonparametric Hidden Semi-Markov Models. *Journal of Machine Learning Research*, 14:673–701, 2013.
- H. Kim, M. Marwah, M. Arlitt, G. Lyon, and J. Han. Unsupervised Disaggregation of Low Frequency Power Measurements. In *Proceedings of the 11th SIAM International Conference on Data Mining*, pages 747–758, Mesa, AZ, USA, 2011.
- J. Z. Kolter and T. Jaakkola. Approximate Inference in Additive Factorial HMMs with Application to Energy Disaggregation. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pages 1472–1482, La Palma, Canary Islands, 2012.
- J. Z. Kolter and M. J. Johnson. REDD: A Public Data Set for Energy Disaggregation Research. In *ACM Knowledge Discovery and Data Mining, workshop on Data Mining Applications in Sustainability*, San Diego, CA, USA, 2011.
- Joon-Ho Lim, Hyunchul Jang, Jaewon Jang, and Soo-Jun Park. Daily activity recognition system for the elderly using pressure sensors. In *30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 5188–5191, August 2008.
- Jesper Møller, Anne Randi Syversveen, and Rasmus Plenge Waagepetersen. Log gaussian cox processes. *Scandinavian journal of statistics*, 25(3):451–482, 1998.
- Shigeru Ohta, Hiroshi Nakamoto, Yoshimitsu Shinagawa, and Tomohiro Tanikawa. A health monitoring system for elderly people living alone. *Journal of telemedicine and telecare*, 8(3):151–156, 2002.
- Oliver Parson, Siddhartha Ghosh, Mark Weal, and Alex Rogers. Non-intrusive load monitoring using prior models of general appliance types. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, pages 356–362, Toronto, ON, Canada, 2012.
- Oliver Parson, Siddhartha Ghosh, Mark Weal, and Alex Rogers. An Unsupervised Training Method for Non-intrusive Appliance Load Monitoring. *Artificial Intelligence*, 217:1–19, 2014.
- P.E. Ross. Managing care through the air [remote health monitoring]. *IEEE Spectrum*, 41(12):26–31, 2004.
- Clíodhna Ní Scanaill, Sheila Carew, Pierre Barralon, Norbert Noury, Declan Lyons, and Gerard M. Lyons. A review of approaches to mobility telemonitoring of the elderly in their living environment. *Annals of Biomedical Engineering*, 34(4):547–563, 2006.
- Hao Song, G Kalogridis, and Zhong Fan. Short paper: Time-dependent power load disaggregation with applications to daily activity monitoring. In *IEEE World Forum on Internet of Things*, pages 183–184, 2014.
- The GPy authors. GPy: A gaussian process framework in python. <http://github.com/SheffieldML/GPy>, 2012–2014.
- T Togawa. Home health monitoring. *Journal of medical and dental sciences*, 45(3):151–160, September 1998.
- Xinpeng Zhang, T. Kato, and T. Matsuyama. Learning a context-aware personal model of appliance usage patterns in smart home. In *IEEE Innovative Smart Grid Technologies - Asia*, pages 73–78, Kuala Lumpur, Malaysia, 2014.
- Mingjun Zhong, Nigel Goddard, and Charles Sutton. Signal aggregate constraints in additive factorial hmms, with application to energy disaggregation. In *Advances in Neural Information Processing Systems*, pages 3590–3598, 2014.
- T Zia, D Bruckner, and A Zaidi. A hidden Markov model based procedure for identifying household electric loads. In *37th IEEE Annual Conference on Industrial Electronics Society*, pages 3218–3223, 2011.
- Jean-Paul Zimmermann, Matt Evans, Jonathan Griggs, Nicola King, Les Harding, Penelope Roberts, and Chris Evans. Household Electricity Survey A study of domestic electrical product usage. Technical report, Intertek, 2012.
- A Zoha, A Gluhak, M Nati, and M A Imran. Low-power appliance monitoring using Factorial Hidden Markov Models. In *Eighth IEEE International Conference on Intelligent Sensors, Sensor Networks and Information Processing*, pages 527–532, 2013.