

Article

Entropy-Based Metrics for Occupancy Detection Using Energy Demand

Denis Hock ^{1,*}, Martin Kappes ¹ and Bogdan Ghita ²

¹ Faculty of Computer Science and Engineering, University of Applied Sciences Frankfurt am Main, 60318 Frankfurt am Main, Germany; kappes@fb2.fra-uas.de

² School of Engineering, Computing and Mathematics, Plymouth University, Plymouth PL4 8AA, UK; bogdan.ghita@plymouth.ac.uk

* Correspondence: dehock@fb2.fra-uas.de

Received: 16 May 2020; Accepted: 29 June 2020; Published: 30 June 2020



Abstract: Smart Meters provide detailed energy consumption data and rich contextual information that can be utilized to assist electricity providers and consumers in understanding and managing energy use. The detection of human activity in residential households is a valuable extension for applications, such as home automation, demand side management, or non-intrusive load monitoring, but it usually requires the installation of dedicated sensors. In this paper, we propose and evaluate two new metrics, namely the sliding window entropy and the interval entropy, inspired by Shannon's entropy in order to obtain information regarding human activity from smart meter readings. We emphasise on the application of the entropy and analyse the effect of input parameters, in order to lay the foundation for future work. We compare our method to other methods, including the Page–Hinkley test and geometric moving average, which have been used for occupancy detection on the same dataset by other authors. Our experimental results, using the power measurements of the publicly available ECO dataset, indicate that the accuracy and area under the curve of our method can keep up with other well-known statistical methods, stressing the practical relevance of our approach.

Keywords: energy demand; entropy applications; privacy

1. Introduction

The smart grid is generally seen as the next-generation power system, connecting all of the actors of the energy system and optimizing the reliability and efficiency of electricity supply and consumption. The advent of ubiquitous IP-based data exchange in modern power grids and the availability of residential energy measurements in low-voltage networks boosted the interest towards appropriate data assessment for time-dependent market prices, non-intrusive load monitoring, or demand side management.

Many of these methods ultimately aim to save energy by optimizing the customers consumption behaviour. Other authors [1,2] provided evidence that buildings are responsible for a major part of the global energy consumption. According to Vafeiadis et al. [3], the detection of occupancy provides the ability to apply several automation applications that may contribute to the reduction of energy consumption. Candanedo et al. [4] found that an accurate determination of occupancy in buildings can save energy, up to 40%.

In this paper, we propose an effective, entropy-inspired profiling mechanism that consolidates the temporal distribution of energy consumption in order to implement an indicator for occupancy. The occupancy is an interesting metric for many new applications in the smart grid, e.g., monitoring and forecast, dynamic prices, or typical load classification. Detecting when a residential household

is occupied, is still a challenging process, which is typically solved by the intrusive deployment of dedicated sensors, such as infrared sensors, magnetic switches, or cameras. Other authors [5,6] generally agree that these setups require the expensive deployment, calibration, and maintenance of hardware. Furthermore, a wide range of machine learning algorithms, which can indicate occupancy using the energy demand, are unsuitable due to the requirement of historic data and high resolution data. Our contribution is that we reveal this information without a priori information, which allows its use without expensive and laborious training. Furthermore, our algorithm can also cope with low resolution measurements.

As human activities that take place without appliances are not reflected in energy demand, we define the term occupancy in this paper as an interaction with electrical appliances or activity of devices. Obviously, mispredictions of occupancy in this sense are sometimes possible as, e.g., a cleverly randomized use of time-switches cannot be distinguished from a human user doing the same activity. While approaches that are aimed at detecting or predicting actual human presence in homes might appear more powerful, we argue that our proposed approach is more efficient as measuring, modelling, and predicting actual human activities, which is either intrusive or difficult. The differences between these approaches are, in fact, negligible from a practical point of view when considering the use of such data for non-intrusive load monitoring, anomaly detection, and energy forecasts.

The idea to use the entropy to analyse the regularity of energy demand is intuitive, but it contains many hidden challenges. We showcase methods to use the energy demand, horizontally with a sliding window and vertically with intervals, as input data for the entropy. We analyse the respective parameters, such as the window size and number of intervals, and interpret the meaning of a high and low entropy. Our aim is to create a comprehensible metric that can be used to identify or distinguish the behaviour of residential households better than raw data.

In the remainder of this paper, we first present an overview of the statistical entropy and discuss details of two independent metrics. We highlight their advantages and limitations regarding their conclusiveness with respect to residential energy demand. Next, we evaluate our metrics by presenting several practical experiments, which illustrate their functionality and features. Our experimental results, which were obtained by applying entropy to real world smart meter data of the public available Electricity Consumption & Occupancy (ECO) dataset, show that the proposed metric can indeed detect occupancy. A comparison of our results with other well-known statistical methods, followed by a summary and potential future work avenues, concludes this paper.

2. Related Work

Many authors agree that residential energy demand can provide unique insights into household characteristics or individual behaviours. An energy provider can use these information to better monitor the complex consumption behaviour, improve the prediction of certain consumer profiles, and suggest optimal energy pricing for particular groups. In particular, Anderson et al. [7] discussed and examined a great many of literature that suggests that we can use residential energy demand to indicate household characteristics and associate consumption patterns to particular consumer groups.

Previous studies of Molina et al. [8] used density-based clustering and supervised learning to identify private information about consumers. Research conducted by Carroll et al. [9] and Mcloughlin et al. [10] suggests that there is a relationship between load curves and the employment status as well as the presence of children. Kolter et al. [11] analysed the relation between demand and building properties, such as the number of rooms and the building value. Beckel et al. [12] extracted the number of occupants from energy demand. Newing et al. [13] associated energy consumption patterns with particular dwellings, income, and number of children.

Nguyen and Aiello [14] as well as Labeodan et al. [15] provide a comprehensive list of approaches to capture occupancy, which are mostly based on sensor data. Kleiminger et al. [5], who provided the Electricity Consumption & Occupancy (ECO) dataset, implemented a method that was based on supervised machine learning algorithms to detect human presence. Jin et al. [6] performed

occupancy detection on both residential and commercial buildings. Becker et al. [16] computed the occupancy with unsupervised occupancy detection algorithms on three public datasets. Peng et al. [17] used unsupervised and supervised learning machine learning methods to predict the occupancy of office buildings.

The major contributions of this paper include the application of an entropy-inspired metric to profile energy demand. The entropy has been extensively applied in other areas such as health care [18], biodiversity assessment [19], or network monitoring [20].

Some of the authors applied the entropy to topics related to the smart grid, but not occupancy detection. E.g., Iranmanesh et al. [21] used the entropy based mutual information method in order to select relevant input features for energy forecasting. Ruiz et al. [22] used the entropy to measure the volatility of energy price markets and relate them to historical events or climatic factors. Gao et al. [23] used the entropy to assess risks of the electricity market. Hock et al. [24] used the entropy to detect energy theft by comparing multiple households.

In contrast to most conventional methods, our metric does not require large amounts of historical data to build a statistical model and it is relatively simple to compute. To the best of our knowledge, an entropy-inspired metric has not been proposed to analyse energy demand before.

3. Method

In the following, we will briefly outline the rationale behind our approach to use a somewhat complex metric to reliably detect events, such as changes in demand and energy usage pattern. Our objective is to detect electric events pointing to occupancy. Due to appliances with high demand, typical statistical ratios are insufficient for reflecting occupancy, because the amplitude of energy intensive or energy saving appliances do not necessarily translate into different human activities. Hence, we propose using regularity or randomness, in other words temporal changes of demand, as a benchmark for occupancy.

Energy demand is not random—the patterns of human activity, given by the stacked load shapes of individual appliances, are often gathered in a few periods per day and repeat daily according to the consumption habits of the household. We aim to observe these sections, which appear due to a high variability as outlier, through measuring entropy.

The entropy is a convenient way to detect outliers in the regularity of energy demand. The distribution of energy demand can be measured horizontally or vertically. By analysing the distribution of demand over time (x-axis), we may find out that most of the demand appears at a specific time, which is an outlier. By analysing the distribution of measurements (y-axis), we may find out that there are only few measurements with high power, which also is an outlier.

Although entropy originated in thermodynamics, Shannon argued with his application of the entropy to information theoretic problems—as a quantitative and qualitative technique for understanding ‘uncertainty’—that entropy has a deeper meaning. It is well known in coding theory that the entropy of a discrete random variable quantifies the average length of the encoding of the random variable. Moreover, Shannon’s entropy measures the average uncertainty, also referred to as information content. The uncertainty is maximized when the outcomes of the random variable are equally likely, which corresponds to a uniform distribution. As will be shown later, entropy-based measures capturing the entropy of energy-consumption over time are useful in detecting occupancy.

Formally, let $L = \{a_1, a_2, \dots, a_n\}$ be a dataset where n is the cardinality of L . Moreover, let x_1, x_2, \dots, x_n denote the frequency of each element in L in some sequence X of elements from the set and $m = \sum_{i=1}^n x_i$ implies the number of all observations. Subsequently, the entropy $H(X)$ is defined by

$$H(X) = - \sum_{i=1}^n \frac{x_i}{m} \cdot \log_2 \left(\frac{x_i}{m} \right) \quad (1)$$

Please note that $0 \leq H(X) \leq \log_2(n)$ where $H(X) = 0$ is assumed if only one element of L occurs in X (by convention, $0 \cdot \log_2(0) = 0$) and $H(X) = \log_2(n)$ if all elements of L occur in X with the same frequency. As $\frac{x_i}{m}$ is synonymous for the occurrence probability, we can extend the interpretation to the statistical distribution of the underlying data: an entropy that is close to $\log_2(n)$ reads as ‘random’, because the elements appear almost equally often, whereas an entropy that is close to zero reads as ‘skewed’, because few elements appear more frequent.

As we use the entropy not in a traditional sense, but as a metric to mathematically trace outliers in a distribution of measurements, the practical function as a metric, which encodes distributions without losing information on outliers, is the most important aspect for us.

Figure 1 illustrates the entropy (top) for a corresponding distribution (bottom) with a generated example. The two figures at the bottom show bar plots with 10 different bars, called (a–j), each visualising a distributions. The distributions are from left to right more and more skewed to showcase the resulting entropy, which decreases with the skewness. The left-hand side focuses on multiple outliers, while the right-hand side focuses on single outliers.

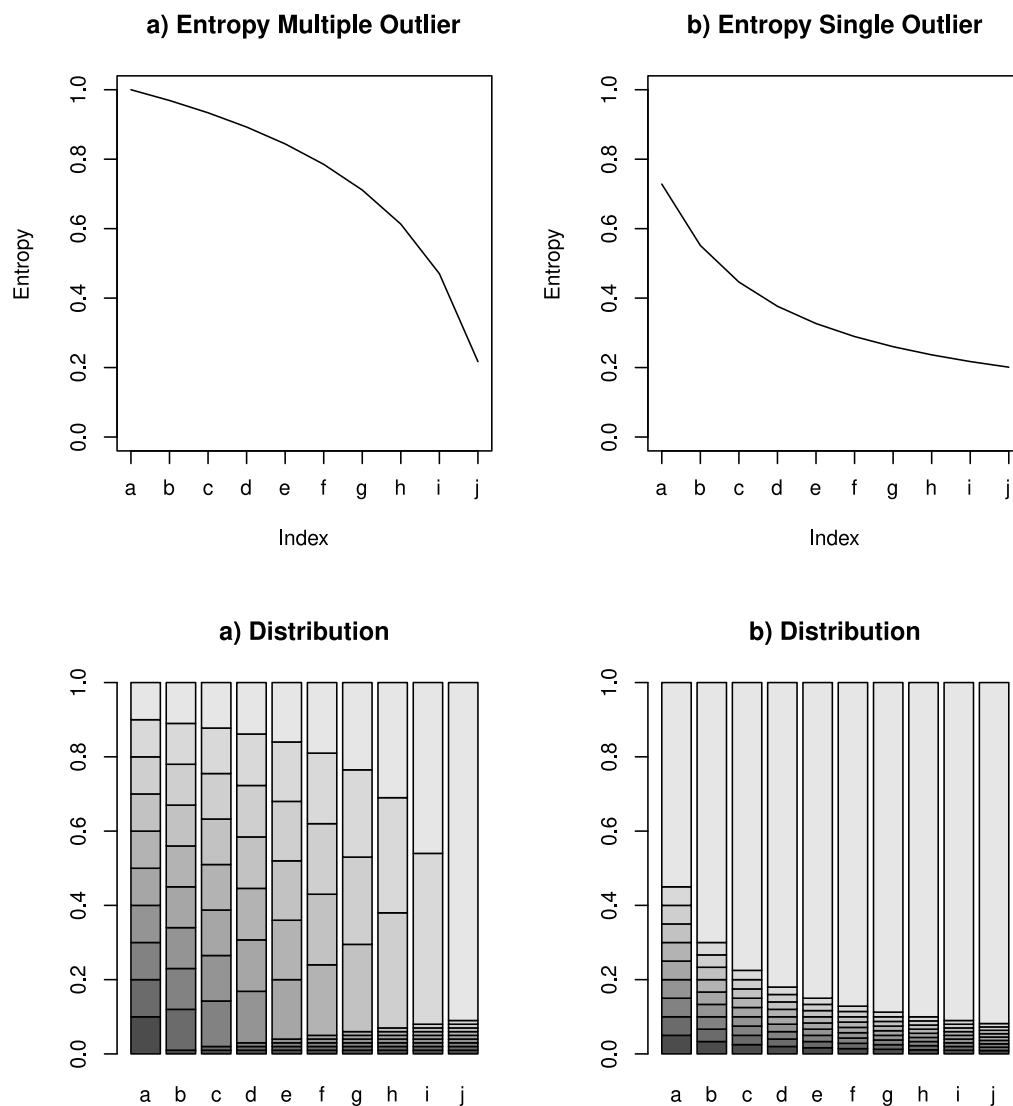


Figure 1. Entropy for uniform and heavily skewed distributions.

Apart from the entropy-inspired metric, other distance functions may also work in order to detect outlier in the distribution. Our metrics result looks strongly correlated to the standard deviation (σ). However, the σ of a vector increases with the distance between values, while the entropy only considers the number of unique values. E.g., for a vector $v_1 = (1W, 1W, 1000W, 1000W)$ and $v_2 = (1W, 250W, 750W, 1000W)$. The σ for v_1 is higher than for v_2 . However, the entropy for v_2 is higher than for v_1 .

4. Application

Subsequently, we illustrate metrics that are based on Shannon's entropy, as a unified analytic approach to discover changes of energy demand over time. As previously explained, the range of the entropy is profoundly affected by the size of the dataset. In our case, it coincides with the finite amount of energy demands n , which, in turn influence its range $[0, \log_2(n)]$. Hence, we will discretise continuous sets, such as \mathbb{R} , in order to use them as input parameter for our entropy-measure.

Following the reasoning of the previous chapter, we introduce two independent methods to observe the uniformity of energy demand vertically and horizontally resulting in two opposite interpretations, different from the information theoretic interpretation, aiming to significantly mirror the changes in energy demand pattern. While our first method sliding window entropy employs the probability to see demand during a certain time, the second method interval entropy conducts the probability to see demand in a certain range. We aim to explain the different interpretation of the result, in contrast to the traditional Shannon entropy, and choose the best approach to detect occupancy.

Human activity stands out due to energy demand peaks. The sliding window entropy can detect such a peak if a single time window contains most of the total demand. Alternatively, human activity stands out due to high variance. The interval entropy can detect this variance if the power measurements are equally distributed over the measurement range.

4.1. Sliding Window Entropy

To summarize the sliding window entropy, we utilize the proportion of energy demand in a time window as input for the entropy to obtain a metric that explains the distribution of energy demand over one day. In contrast to the information theoretic entropy, our metric is maximized if the demand in all of the time windows is equal and the consumption is distributed uniformly throughout the day and minimized when the consumption is concentrated in a single interval. Hence, we assume occupancy when the entropy is minimized.

Consider a finite time series $\mathcal{T} = a_1, a_2 \dots a_n$, $a_i \in \mathbb{R}_0^+ \forall 1 \leq i \leq n$, with n elements - representing energy demand. Subsequently, a sliding window \mathcal{X} with size m , where the sliding step ($1 \leq \Delta \leq m$) defines the number of elements by which the window slides each iteration t , can be defined as $\mathcal{X}_t = a_{t\Delta}, a_{t\Delta+1}, \dots, a_{t\Delta+m-1}$, which results in a total of $0 \leq t \leq \frac{n-m}{\Delta} + 1$ windows to cover all elements. The traditional entropy considers the occurrences x_i over the data length n , which is equivalent to the probability to see a certain element $\frac{x_i}{n}$. In our energy usage context, the entropy considers the demand probability in a sliding window $P(\mathcal{X}_t)$, which is equivalent to the normalised integrated area under the demand line. We simplify the approximation of area, by dividing the demand of a window $h(\mathcal{X}_t)$, which is the sum of all measurements in a window, by the total demand $h(\mathcal{T})$:

$$P(\mathcal{X}_t) = \frac{h(\mathcal{X}_t)}{h(\mathcal{T})} \quad (2)$$

We could also interpret the probability as the proportion of daily energy used during that sliding window. Note that the sum of all probabilities can be greater than 1, if the sliding windows overlap ($\Delta < m$), which results in a scaled entropy.

Figure 2 illustrates the energy demand (grey line) divided into three time window $\mathcal{X}(1 : 3)$, $m = 8h$, $\Delta = m$ (dotted lines), with the probability for energy demand delineated in black.

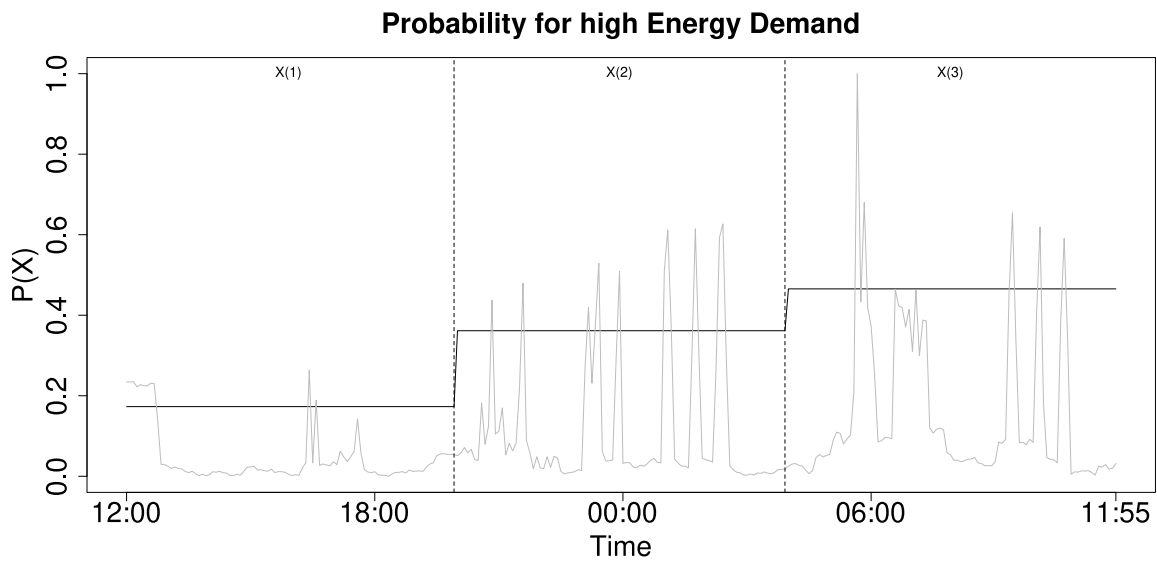


Figure 2. Integrated demand for three time windows.

The entropy is maximized if the demand is distributed uniformly and minimized if the total demand is only present in a single sliding window. The result ranges from $[0, \log(\frac{n-m}{\Delta} + 1)]$. Equation (3) illustrates how to calculate the entropy while using the above probability.

$$H(\mathcal{X}) = - \sum_{t=0}^{\frac{n-m}{\Delta} + 1} P(\mathcal{X}_t) \cdot \log(P(\mathcal{X}_t)) \tag{3}$$

Occupancy stands out due to the imbalance and change of energy consumption, which manifests here in form of demand peaks. High demand concentrated in a single time window results in a minimized sliding window entropy. Figure 3 visualizes the complete process to derive the occupancy from raw energy demand. We first compute the probability for high demand (middle) in three different resolutions and then compute the entropy for a time window of three hours using the probability as input data. In the first step, we define three different sliding windows \mathcal{X} with $m_1 = 15$, $m_2 = 30$, and $m_3 = 60$ each with $\Delta = m$ over the total of $n = 1440$ measurements. Step 2 shows the resulting 24 probabilities for m_1 (dot), 48 probabilities for m_2 (triangle) and 96 probabilities for m_3 (cross). Here, the sum of measurements in each sliding window is divided by the sum of the complete time series, in order to generate one probability value for each sliding window. In the last step, the probabilities derived from the previous time windows $P(\mathcal{X}_t)$ are used as input value to compute an entropy. In our example, we use three probabilities (dot), six probabilities (triangle), and 12 probabilities (cross), which always corresponds to three hours, as input to compute the entropy. This results in eight entropy values per day, but the amount of probabilities used to compute the entropy is of course a variable parameter.

In the resulting plot, we can clearly see the influence of several parameters, namely the number of time windows used as input for the entropy (middle) and the size of the time window for each entropy value (bottom). The number of probabilities changes the level of our output—the resulting entropy is smaller the more input values we have, while the pattern of our curve is unchanged. The entropy results in a filter operation over the energy demand and smaller time windows for the entropy would show more detail. There is a visible drop of the entropy whenever the amplitude significantly changes within a time window. For our decision, whether a time window represents occupancy or not, we can simply utilize a threshold. What may look unusual is that the difference between time windows is not considered in this metric, because the sum of probabilities of any time window result in 100%. Hence, two time windows with $(0.1W, 1W)$ and $(1W, 100W)$ result in the same entropy.




Figure 3. Probability to use energy in a certain time window.

4.2. Interval Entropy

To summarize the interval entropy, we count the amount of measurements in each interval to compute a probability that results in a metric affected by the distribution between intervals. This metric is maximized if the demand values are uniformly distributed over all intervals, which is the direct opposite to our previously introduced sliding window entropy.

Consider a finite time series $\mathcal{T} = a_1, a_2 \dots a_n$, $a_i \in \mathbb{R}_0^+ \forall 1 \leq i \leq n$, with n elements. Subsequently, a finite number m of equal-sized discrete intervals \mathcal{I} covering a total range of $r = \max(\mathcal{T}) - \min(\mathcal{T})$, is given by $\mathcal{I}_i = [\frac{r}{m} \cdot (i - 1) + \min(\mathcal{T}), \frac{r}{m} \cdot i + \min(\mathcal{T})]$, with $1 \leq i \leq m$.

Our entropy in Equation (4) represents the probability that the demand is an element from a certain interval \mathcal{I} , where n is the number of elements in \mathcal{T} and m is the number of intervals.

$$\begin{aligned}
 \text{Freq}(\mathcal{I}_i) &= \sum_{j=0}^n a_j \in \mathcal{I}_i \\
 P(\mathcal{I}_i) &= \frac{\text{Freq}(\mathcal{I}_i)}{n}
 \end{aligned}
 \tag{4}$$

Equation (4) shows the computation of the input for the entropy. For each interval \mathcal{I}_i , we count the number of measurements a using the function $\text{Freq}()$ and then divide by the total number of measurements to get the probability. The entropy is maximized if the demand distributes an equal number of elements to each interval m and minimized if the demand involves only one interval. The result ranges between $[0, \log(m)]$. Equation (5) show how to utilize the probability of each interval to compute the entropy.

$$H(\mathcal{I}) = - \sum_{i=0}^m P(\mathcal{I}_i) \cdot \log(P(\mathcal{I}_i))
 \tag{5}$$

We experimented with different methods to generate intervals (see Figure 4)—the top figures show equally distributed intervals, while the bottom figures have intervals with dynamic size—namely, logarithmically distributed and clustered to high variance sections—which enables us to also consider small changes of energy demand. Note that this figure shows the raw data. In order to compute the entropy, we count the number of measurements in each interval and compute the probability by dividing through the total number of intervals. Next, we compute the entropy using the probabilities as input data.




Figure 4. Energy demand in different intervals.

We found that large intervals react well on changing variance and reliably show the start and end of each phase, while many small intervals may emphasize too much on small changes. The clustering approach might be difficult to implement in real world scenarios, since the clusters would differ for each household and the distribution of cluster centres may show legitimate changes each month.

Overall, our method assumes human activity to fall in a certain range of the entropy and, hence, defines a threshold to split the metric. The binary result $r \in \{0, 1\}$ depends on a threshold ε .

$$r = \begin{cases} 1 & \text{if } H \text{ is } > \varepsilon \\ 0 & \text{if } H \text{ is } \leq \varepsilon \end{cases} \quad (6)$$

Note that the interpretation of the sliding window entropy and interval entropy is oppositional and, hence, the meaning of r changes. If the sliding window entropy is higher ε , then the energy demand is equally distributed over many sliding windows and, hence, does not contain the peaks that indicate human activity. If the interval entropy is higher ε , then the measurements are distributed over all intervals and, hence, have a high variation, which indicates human activity.

In summary, the sliding window entropy depends upon the size and overlap of the window, while the interval entropy depends upon the number of intervals and distribution of intervals. After computing the probabilities with these parameters, both of the methods can be further tuned with the number of probabilities used as input. As the values in our metric are always relative to each other, less input values mean that an outlier that leads to a low entropy is usually more visible. The parameter tuning can be seen as an optimization problem. In the next section, we graphically analyse the influence of the parameters, but a detailed method to automatize this process for different households is beyond the scope of this paper.

5. Results

In the following experiments, we aim to detect occupancy while using real world measurements and compare our method with other well established prediction schemes. We first introduce our method to create a ground truth by labelling occupancy using appliance level data. Next, we predict these ground truth data without access to the appliance level data. Note that the ground truth data cannot be reliably created without the appliance level data, which are not available for the algorithms that aim to detect the occupancy.

We use the ECO dataset provided by Kleiminger et al. [5], which is a comprehensive dataset for non-intrusive load monitoring. The dataset offers individual appliance readings every second (Hz). The measurements used here are mean-cycle-power measurements of the real power (W), as provided by smart meters. The data are measured at the plug level for individual appliances and, apart from declared exceptions, provides measurements in Watt with four decimal places. Because we do not attempt to detect individual appliance signatures, we aggregated the values to one measurement every minute, which is due to the privacy concerns a more realistic interval for measurements in real world smart grids.

We used the appliance-level data of all six household for summer (2012 June to 2012 October) and winter (2012 October to 2013 January) in order to create the ground truth data. We labelled events that point to occupancy by selecting appliances with non-periodic behaviour (7c) and exclude appliances in stand-by (7a) mode, appliances running permanently and appliances with regular energy pattern (7b) (e.g., freezer, fridge, ...).

Consider each of the i appliances in the ECO dataset as finite time series $\mathcal{T}_i = \{a_0..a_n \mid a \in \mathbb{R}\}$, with n elements. We select each appliance \mathcal{T}_i that meets following conditions:

$$C_1 = \max(\mathcal{T}_i) > 10 \quad (7a)$$

$$C_2 = Q2(\mathcal{T}_i) - \min(\mathcal{T}_i) < \max(\mathcal{T}_i) - Q2(\mathcal{T}_i) \quad (7b)$$

$$C_3 = Q3(\mathcal{T}_i) - Q2(\mathcal{T}_i) < \max(\mathcal{T}_i) - Q3(\mathcal{T}_i) \quad (7c)$$

Note that, Q2 denotes the second quantile, which is the median, and Q3 is the third quartile which is the middle value between the median and max. If an appliance meets these conditions, than all values greater than the arithmetic mean value of the appliance were labelled as occupancy.

We compare the entropy inspired metric with two methods provided by [16], namely the and (GeoMa).

With the data outlined above, we compare our own method to four well-known algorithms, namely (a) Page–Hinkley test, (b) Geometric Moving Average (c) Moving Average, (d) Seasonal Decomposition, and the (e) Standard Deviation. Each algorithm aims to classify occupancy (true or false) for 15 min time windows.

The first two algorithms are provided by Becker et al. [16], who performed occupancy detection on the same dataset. The method (a) Page–Hinkley Test (PHT) is a change detection algorithm, while (b) the Geometric Moving Average (GeoMA) detects any value higher than the average as occupied. Method (c) detects whether a measurement has a demand higher or lower than the Moving Average (MA), plus an adjustment value, (d) uses the trend component of Seasonal Decomposition (SD) [25] and detects values higher than a threshold. The last method, (e) detects whether the standard deviation is higher than the daily main. Our two original algorithms are evaluated, as follows, the sliding window entropy is using five minutes to compute a probability for the 15 min time window, whereas the interval entropy is using 20 uniformly distributed intervals. Note that we had to accept some limitations due to the different input and output data of each algorithm. Namely, some of the algorithms require a time window as input data, while others generate a decision for each input. For the latter, we aggregated the measurements to 15 min intervals and did not find a decreased performance, which is in consensus with Becker et al. [16]. Furthermore, we usually evaluate our algorithms with the so-called Area under the curve (AUC) of a receiver operating characteristic (ROC) curve in order to avoid to falsify results due to a non-optimal threshold. The curve shows the sensitivity (true positive rate) and the specificity (true negative rate) for any possible threshold.

The sensitivity is given by $\frac{TP}{TP+FN}$, while the specificity is given by $\frac{TN}{TN+FP}$. The accuracy of a method can be defined by the AUC. A good method can maximize both sensitivity and specificity, which results in a big area under the curve, while a random method results in a diagonal line and AUC of 0.5.

However, this is not possible for all algorithms as some of them require the threshold throughout the computation and only return a decision instead of an anomaly score. We still computed the AUC with a boolean result, (e.g., to perform a statistical test), but had to optimize the threshold for each function manually using linear optimization on the summer half of the first household of the ECO dataset.

In Table 1, we computed the AUC by evaluating the 96 results per day (15 min intervals) with the expected ground truth. A good method can maximize both sensitivity and specificity, which results in a big area under the curve, while a random method results in a diagonal line and an AUC of 0.5. Hence, the probability that an algorithm assigns the correct classification can be defined by the area under the curve (AUC). Table 1 shows the performance of each algorithm. The decomposition performed very good on this month, but we were not able to adjust the threshold for a good performance on the complete dataset. It might be possible to increase the performance by dynamically finding a new threshold after a fixed time period. However, this is not in the scope of this work.

Table 1. Results of all methods (AUC).

	PHT	GeoMa	MA	SA	Std. Dev.	Int. Ent.	Win. Ent.
2012-06-01	0.837	0.765	0.648	0.868	0.754	0.897	0.912
2012-06-02	0.511	0.744	0.559	0.67	0.534	0.506	0.744
2012-06-03	0.672	0.758	0.691	0.823	0.711	0.871	0.806
2012-06-04	0.457	0.815	0.672	0.967	0.955	0.803	0.809
2012-06-05	0.747	0.794	0.688	0.824	0.763	0.793	0.831
2012-06-06	0.723	0.875	0.823	0.913	0.779	0.898	0.905
2012-06-07	0.602	0.63	0.672	0.69	0.58	0.637	0.598
2012-06-08	0.845	0.802	0.746	0.893	0.81	0.942	0.885
2012-06-09	0.477	0.744	0.749	0.824	0.861	0.917	0.718
2012-06-10	0.853	0.67	0.612	0.896	0.641	0.702	0.79
2012-06-11	0.836	0.664	0.793	0.867	0.745	0.9	0.817
2012-06-12	0.966	0.881	0.581	0.874	0.784	0.897	0.924
2012-06-13	0.618	0.821	0.669	0.875	0.685	0.779	0.769
2012-06-14	0.707	0.692	0.668	0.832	0.784	0.879	0.778
2012-06-15	0.416	0.649	0.417	0.409	0.549	0.504	0.513
2012-06-16	0.8	0.908	0.748	0.922	0.829	0.95	0.975
2012-06-17	0.824	0.548	0.499	0.5	0.952	0.537	0.702
2012-06-18	0.831	0.789	0.682	0.923	0.717	0.856	0.885
2012-06-19	0.716	0.846	0.677	0.9	0.752	0.837	0.91
2012-06-20	0.699	0.793	0.622	0.896	0.637	0.692	0.919
2012-06-21	0.735	0.656	0.579	0.712	0.588	0.626	0.756
2012-06-22	0.592	0.626	0.523	0.532	0.516	0.51	0.479
2012-06-23	0.79	0.763	0.693	0.943	0.823	0.868	0.769
Total (Avg.)	0.707	0.749	0.653	0.807	0.728	0.774	0.791

Figure 5 shows the results on a larger time scale and for all households. Each bar visualizes the result of an algorithm in summer and winter. The results are given in Accuracy (ACC), which is computed with a confusion matrix, given by $(tp + tn)/(tp + tn + fp + fn)$, which returns the performance in percentage.

Next, we perform a pairwise bootstrap test (see Table 2) with the results of two methods, in order to evaluate the significance of the AUC over the summer in household 1. The bootstrap test re-samples the original results repeatedly ($N = 2000$), which approximately follows a normal distribution, in order to perform a hypothesis test. Here, the null hypothesis is that both methods perform equally and the alternative hypothesis is that the AUC of both methods are different. A small p -value shows that the null hypothesis can be rejected.

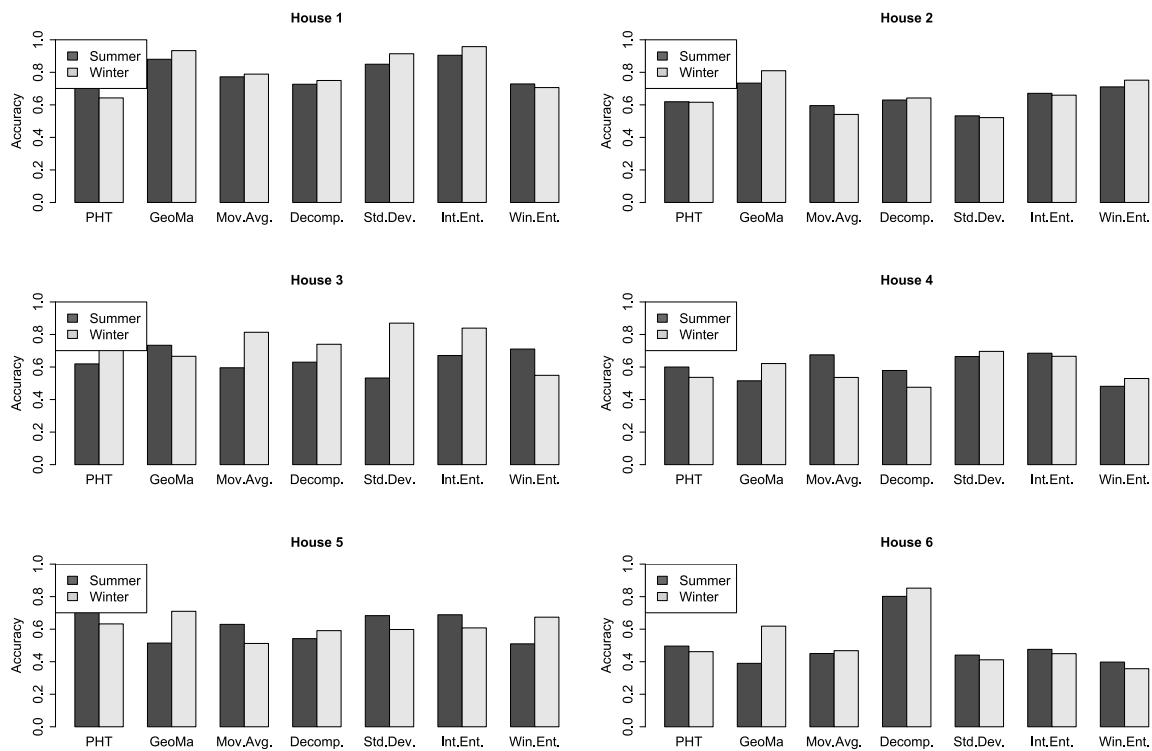


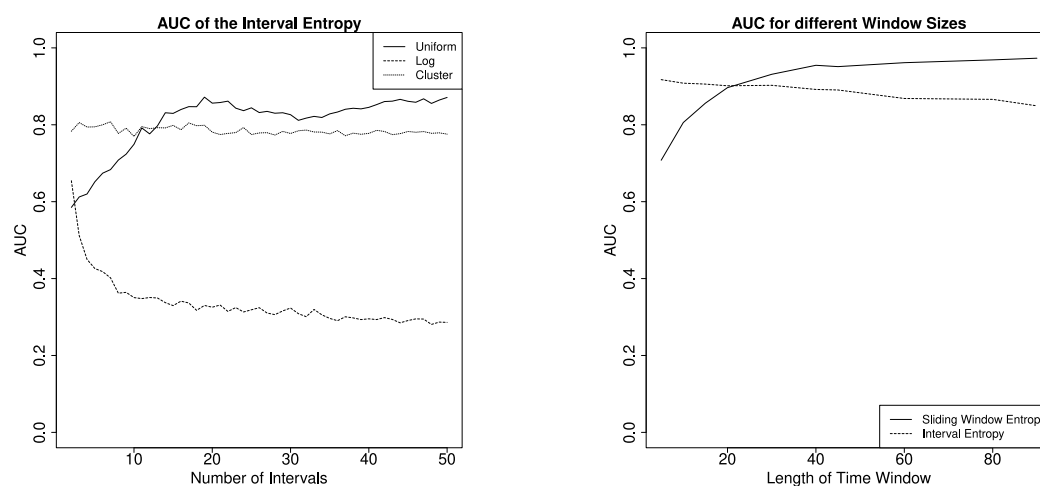
Figure 5. Accuracy for each household of the ECO dataset.

Table 2. Bootstrap test for all methods (AUC).

H_1	AUC1	AUC2	p -Value
PHT vs. GeoMA	0.718	0.779	5.28×10^{-26}
PHT vs. MA	0.718	0.639	8.83×10^{-46}
PHT vs. SD	0.718	0.661	3.28×10^{-21}
PHT vs. Std. Dev.	0.718	0.722	5.70×10^{-1}
PHT vs. Int. Ent.	0.718	0.745	2.15×10^{-5}
PHT vs. Win. Ent.	0.718	0.714	5.61×10^{-1}
GeoMA vs. MA	0.779	0.639	2.24×10^{-136}
GeoMA vs. SD	0.779	0.661	1.88×10^{-70}
GeoMA vs. Std. Dev.	0.779	0.722	2.84×10^{-20}
GeoMA vs. Int. Ent.	0.779	0.745	9.83×10^{-11}
GeoMA vs. Win. Ent.	0.779	0.714	3.60×10^{-29}
MA vs. SD	0.639	0.661	7.30×10^{-4}
MA vs. Std. Dev.	0.639	0.722	8.44×10^{-33}
MA vs. Int. Ent.	0.639	0.745	3.45×10^{-58}
MA vs. Win. Ent.	0.639	0.714	3.87×10^{-29}
SD vs. Std. Dev.	0.661	0.722	1.59×10^{-16}
SD vs. Int. Ent.	0.661	0.745	7.48×10^{-33}
SD vs. Win. Ent.	0.661	0.714	4.40×10^{-11}
Std. Dev. vs. Int. Ent.	0.722	0.745	1.12×10^{-8}
Std. Dev. vs. Win. Ent.	0.722	0.714	2.65×10^{-1}
Int. Ent. vs. Win. Ent.	0.745	0.714	4.88×10^{-6}

Last but not least, we evaluated the parameters for the interval entropy and sliding window entropy. Here, we used the same input parameters as above to compute the AUC, but for different time windows and in the case of the interval entropy also for different numbers and types of intervals (note that, a time window of 60 Minutes means that we compute the entropy 24 times per day).

In Figure 6, we can see the influence of parameters on the AUC of the interval entropy in detail. The output varies depending on the number of intervals (left plot), where each line shows a different type to distribute the intervals. Here, the uniform distribution showed the best results and it is maximized with about 20 intervals and time windows of 30 min. Next, we compare the AUC of the interval entropy (20 uniform intervals) and sliding window entropy with different window sizes. The sliding window entropy performs better with larger time windows, while the AUC of the interval entropy is slowly decreasing with greater window sizes. The interval entropy works on time periods of arbitrary length. Although using it on long periods is possible (e.g., to detect absence on vacation or incoming tenants), we did not test the metric for such parameters. We think that it is more desirable to have an higher output resolution and repeat the process several times until the complete data is covered. When adjusting parameters, such as the length of input and output, for both the interval and sliding window entropy, it is intuitive to think of entropy as the probability that all input values are similar. If more probabilities are used as input values, we have more similar values and the entropy is generally higher. Hence, the difference between a normal and anomalous entropy is smaller, which makes it challenging to define a good threshold. The parameter tuning obviously depends on the actual household and slightly changes every time. However, we can clearly see that especially the sliding window entropy does not depend on fine granular data, as one measurement per time window is theoretically enough for this method to work.



(a) Different intervals with a window size of 15 min.

(b) Different window sizes with 20 intervals.

Figure 6. Effect of parameters on the interval entropy.

6. Conclusions and Outlook

In this paper, we presented a non-intrusive approach to detect occupancy without a priori data. Our first results compared the feasibility of an entropy inspired metric while using different parameter. Our results showed a significantly reduced noise, outlier resilience and clearly visible usage patterns. Both of the methods reflect changes in energy demand with slightly different purpose. The sliding window entropy accurately shows on-/off behaviour and peaks of appliances while hiding less interesting absolute values. The interval entropy highlights high variance sections of a time-series while masking static behaviour. Both of the methods appear to behave well with a moderate number of input values. Our experiments utilizing the ECO dataset motivated the real world suitability.

In the graphical representation of the accuracy for all households, the results vary for all households as well as winter and summer. The differences for winter and summer were very different for each household, only 'house 3' performed clearly better during winter. Altogether, we received results varying from 0.7 to above 0.9. For 'house 1', which used the set of the parameters for all methods, we reached an accuracy higher 0.9. This is an indicator that future approaches can improve the results by tuning the parameter for a certain household, distinguish seasons, and perhaps distinguish weekend

and weekdays. In our experiments, the PHT and SA also performed good and similar to the entropy look at the randomness of the energy demand. The interval entropy performed better than the sliding window entropy, but the parameter tuning was more difficult, and we experienced more noise in the performance depending on the type and number of intervals.

In the experimental evaluation, the entropy could not clearly exceed the results of all other methods. However, we had to accept many restrictions in order to compare our methods to other methods (e.g., the resolution, window sizes, and required boolean output) and think it is possible to further increase the performance by fine-tuning some parameters. Here, our focus was set on the entropy as a metric, which is intuitive to interpret and valuable as a numeric value to analyse and understand energy demand. We aimed to present the fundamental knowledge necessary to apply this new metric to other applications of energy demand, such as energy theft detection and non-intrusive load monitoring.

We believe that the ‘information content’ is profoundly affected by many energy related events and, hence, the entropy should be generally suitable for an eclectic range of classification and identification tasks when combined with appropriate parameters. As future prospects, we would like to combine our metric with tasks, such as non-intrusive load monitoring, anomaly detection, and energy demand forecasting.

Author Contributions: Conceptualization, D.H.; Data curation, D.H.; Formal analysis, D.H.; Funding acquisition, M.K.; Investigation, D.H.; Methodology, D.H.; Project administration, M.K.; Resources, D.H.; Software, D.H.; Supervision, M.K. and B.G.; Validation, D.H., M.K. and B.G.; Visualization, D.H.; Writing—original draft, D.H.; Writing—review & editing, M.K. and B.G. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the German Ministry of Economics and Technology under the Central Innovation Programme for SMEs (ZIM) Grant No. ZF4131804HB8.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

H	Entropy
σ	Standard Deviation (Std. Dev.)
Δ	Sliding Step
\mathcal{T}	Time Series
\mathcal{X}	Sliding Window
\mathcal{I}	Interval
\mathbb{R}	Real Numbers
Hz	Hertz (unit frequency)
W	Watt (unit power)
P	Probability
TP	True Positive
TN	True Negative
FP	False Positive
FN	False Negative
PHT	Page-Hinkley Test
AUC	Area under the curve
ROC	Receiver Operating Characteristic
ECO	Electricity Consumption & Occupancy
GeoMA	Geometric Moving Average
MA	Moving Average
SA	Seasonal Decomposition

References

1. Ahmad, J.; Larijani, H.; Emmanuel, R.; Mannion, M.; Javed, A. Occupancy detection in non-residential buildings—A survey and novel privacy preserved occupancy monitoring solution. *Appl. Comput. Inform.* **2018**, *14*, 107–224. [[CrossRef](#)]
2. Passer, A.; Ouellet-Plamondon, C.; Kenneally, P.; John, V.; Habert, G. The impact of future scenarios on building refurbishment strategies towards plus energy buildings. *Energy Build.* **2016**, *124*, 153–163. [[CrossRef](#)]
3. Vafeiadis, T.; Zikos, S.; Stavropoulos, G.; Ioannidis, D.; Krinidis, S.; Tzovaras, D.; Moustakas, K. Machine learning based occupancy detection via the use of smart meters. In Proceedings of the 1st International Symposium on Computer Science and Intelligent Controls (ISCSIC), Budapest, Hungary, 20–22 October 2017; pp. 6–12.
4. Candanedo, L.M.; Feldheim, V. Accurate occupancy detection of an office room from light, temperature, humidity and CO₂ measurements using statistical learning models. *Energy Build.* **2016**, *112*, 28–39. [[CrossRef](#)]
5. Kleiminger, W.; Beckel, C.; Santini, S. Household occupancy monitoring using electricity meters. In Proceedings of the 3rd International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp), Osaka, Japan, 9–11 September 2015; pp. 975–986.
6. Jin, M.; Jia, R.; Spanos, C.J. Virtual occupancy sensing: Using smart meters to indicate your presence. *IEEE Trans. Mob. Comput.* **2017**, *16*, 3264–3277. [[CrossRef](#)]
7. Anderson, B.; Lin, S.; Newing, A.; Bahaj, A.; James, P. Electricity consumption and household characteristics: Implications for census-taking in a smart metered future. *Comput. Environ. Urban Syst.* **2017**, *63*, 58–67. [[CrossRef](#)]
8. Molina-Markham, A.; Shenoy, P.; Fu, K.; Cecchet, E.; Irwin, D. Private memoirs of a smart meter. In Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building, Zurich, Switzerland, 2 November 2010; pp. 61–66.
9. Carroll, J.; Lyons, S.; Denny, E. Reducing household electricity demand through smart metering: The role of improved information about energy saving. *Energy Econ.* **2014**, *45*, 234–243. [[CrossRef](#)]
10. McLoughlin, F. Characterising Domestic Electricity Demand for Customer Load Profile Segmentation. Ph.D. Thesis, Dublin Institute of Technology, Dublin, Ireland, 2013.
11. Kolter, J.Z.; Ferreira, J., Jr. A large-scale study on predicting and contextualizing building energy usage. In Proceedings of the 25th Conference on Artificial Intelligence (AAAI), San Francisco, CA, USA, 7–11 August 2011.
12. Beckel, C.; Sadamori, L.; Santini, S. Automatic socio-economic classification of households using electricity consumption data. In Proceedings of the 4th International Conference on Future Energy Systems (e-Energy), Berkeley, CA, USA, 22–24 May 2013; pp. 75–86.
13. Newing, A.; Anderson, B.; Bahaj, A.; James, P. The role of digital trace data in supporting the collection of population statistics—the case for smart metered electricity consumption data. *Popul. Space Place* **2016**, *22*, 849–863. [[CrossRef](#)]
14. Nguyen, T.A.; Aiello, M. Energy intelligent buildings based on user activity: A Survey. *Energy Build.* **2013**, *56*, 244–257. [[CrossRef](#)]
15. Labeodan, T.; Zeiler, W.; Boxem, G.; Zhao, Y. Occupancy measurement in commercial office buildings for demand-driven control applications—A survey and detection system evaluation. *Energy Build.* **2015**, *93*, 303–314. [[CrossRef](#)]
16. Becker, V.; Kleiminger, W. Exploring zero-training algorithms for occupancy detection based on smart meter measurements. *Comput.-Sci.-Res. Dev.* **2018**, *33*, 25–36. [[CrossRef](#)]
17. Peng, Y.; Rysanek, A.; Nagy, Z.; Schlüter, A. Using machine learning techniques for occupancy-prediction-based cooling control in office buildings. *Appl. Energy* **2018**, *211*, 1343–1358. [[CrossRef](#)]
18. Yan, X.; Zhang, L.; Li, J.; Du, D.; Hou, F. Entropy-based measures of hypnopompic heart rate variability contribute to the automatic prediction of cardiovascular events. *Entropy* **2020**, *22*, 241. [[CrossRef](#)]
19. Rahimian Boogar, A.; Salehi, H.; Pourghasemi, H.R.; Blaschke, T. Predicting habitat suitability and conserving juniperus spp. habitat using SVM and maximum entropy machine learning techniques. *Water* **2019**, *11*, 2049. [[CrossRef](#)]
20. Kawan, C. Entropy in networked control. *Entropy* **2019**, *21*, 392. [[CrossRef](#)]

21. Iranmanesh, H.; Abdollahzade, M.; Miranian, A. Mid-term energy demand forecasting by hybrid neuro-fuzzy models. *Energies* **2012**, *5*, 1–21. [[CrossRef](#)]
22. Ruiz, M.d.C.; Guillamón, A.; Gabaldón, A. A new approach to measure volatility in energy markets. *Entropy* **2012**, *14*, 74–91. [[CrossRef](#)]
23. Gao, B.; Wu, C.; Wu, Y.; Tang, Y. Expected Utility and Entropy-Based Decision-Making Model for Large Consumers in the Smart Grid. *Entropy* **2015**, *17*, 6560–6575. [[CrossRef](#)]
24. Hock, D.; Kappes, M.; Ghita, B. Using multiple data sources to detect manipulated electricity meter by an entropy-inspired metric. *Sustain. Energy Grids Netw.* **2020**, *21*, 100290. [[CrossRef](#)]
25. Kendall, M.; Stuart, A. *The Advanced Theory of Statistics*; Charles Griffin and Co., Ltd.: London, UK, 1983; Volume 3.



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).