

CONCEPTUAL FRAMEWORK FOR ENHANCING ENERGY
DISAGGREGATION USING CROWD AND CONTEXT

by

AHMED RAZA SAGARWALA

B.Tech, Ryerson University, 2008

A Major Research Paper

Presented to Ryerson University

In partial fulfillment of the
Requirements for the degree of
Master of Digital Media

In the Yeates School of Graduate Studies

Toronto, Ontario, Canada, 2015

© Ahmed Sagarwala, 2015

Author Declaration

AUTHOR'S DECLARATION FOR ELECTRONIC SUBMISSION OF A MRP

I hereby declare that I am the sole author of this MRP. This is a true copy of the MRP, including any required final revisions.

I authorize Ryerson University to lend this MRP to other institutions or individuals for the purpose of scholarly research.

I further authorize Ryerson University to reproduce this MRP by photocopying or by other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research.

I understand that my MRP may be made electronically available to the public.

Signed

Ahmed Sagarwala

Abstract

This paper explores existing electrical disaggregation workflows and how they can be augmented with context awareness through datasets. The goal of energy disaggregation is to educate consumers on their energy usage. Additional benefits in automation, security, and energy auditing can be realized through disaggregation. The use of statistical analysis provides specific device consumption information that can be actioned to conserve energy in a directed and methodical manner. The current landscape of disaggregation is a complex workflow involving algorithms that detect, analyze and reveal consumption patterns. Disaggregation workflows involve the acquisition of energy signals for an entire building, refining readings, detecting events, extracting features, and classification. Each step in the workflow impacts the accuracy in which individual devices are detected. Disaggregation workflows may incorporate device usage and weather patterns to improve accuracy, but crowdsourcing signatures and the incorporation of datasets that allow for context awareness are strategies yet to be adopted.

Acknowledgements

I would like to express my appreciation to the people that have made my experience at Ryerson University inspiring and challenging. It's the people that see issues and realize they have the ability to influence change at a fundamental level. Those that see opportunities can be created in all situations. I want to recognize everyone involved in the Faculty of Communication and Design and Yeates School of Graduate Studies. Especially program administrators like Sonya Tacconne and Marietta Canlas who work very hard at keeping day-to-day operations on track.

For the opportunities I've received to teach and research, I would like to thank Matthew Kyan (my supervisor), Richard Adams, Ian Baitz, Jaigris Hodson, Omar Grant, and Dr. Ira Levine. To my second reader, Michael Carter, for taking the time to meet with me and review my major research paper.

Writing this major research paper has left my house without power and funds several times while I took apart my breaker panel and installed monitoring hardware. This allowed me to aid my understanding of some of the concepts covered in the paper. My family has put up with less income and many power and network outages so that I could pursue my goals in education. To make things better, I would like to acknowledge Ryerson's Dennis Mock Scholarship and miscellaneous financial and mentoring assistances received through the LaunchZone, Masters in Digital Media, and Digital Media Experience at Ryerson.

The equipment to conduct my research has been developed by The Energy Detective, Efergy, Blueline, and Energate. Software by Adobe, Bidgely, Plottwatt, TED, and NILM-TK. Powerstream (my utility provider in Vaughan, Ontario) allowed me to participate in beta studies that gave me a new perspective on the distribution and tracking of electrical consumption. I especially feel it important to acknowledge the open-source community at large that has created so many tools and made them available to the public.

I dedicate this paper to the many people forgetting to turn off their devices.

Table of contents

Author Declaration	ii
Abstract.....	iii
Acknowledgements.....	iv
Table of contents	v
List of tables	vii
List of figures.....	viii
Lists of symbols, abbreviations or other	ix
1. Background & Rationale.....	1
2. Approaches to Electrical Monitoring	3
2.1 Bills & Statements	4
2.2 Plugin Metres	4
2.3. Smart Metres	6
2.4. Induction Clamp	9
3. Opportunities with Disaggregation	11
4. The Building Blocks of Disaggregation	12
4.1. Power Draw	15
4.2. States.....	16
4.3. Frequency of Use	17
5. Algorithms for Disaggregation	17
5.1. Factorial Hidden Markov Model (FHMM).....	17
5.2. Artificial Neural Networks (ANN)	19

5.3. Support Vector Machines (SVM).....	20
5.4. Nearest Neighbour (k-NN)	20
5.5. Dynamic Time Warping (DTW).....	20
6. Disaggregation Workflow	21
6.1. Measure Aggregate	21
6.2. Store Aggregate Reading.....	21
6.3. Initial Pass	21
6.4. Data Segmentation.....	22
6.5. Signature Matching.....	22
7. Furthering the Workflow	23
7.1. Workflow Using Crowd-sourcing.....	26
8. Integrating External Datasets.....	28
8.1. User Habits and Trends	28
8.2. Conceptual Workflows for Context Awareness.....	31
9. Conclusion	41
9.1. Summary of Findings	41
9.2. Future Considerations/Research	42
10. Bibliography of Sources & Materials	44

List of tables

Table 1: Comparison of disaggregation benefits by group	2
Table 2: Comparison of tracking methods	4
Table 3: Location awareness technologies	33
Table 4: Comparison of location services by median error	36

List of figures

Figure 1: Smart metre penetration by region	7
Figure 2: Breaker box with 2 phase configuration.....	10
Figure 3: Induction clamp installation	10
Figure 4: Aggregate and disaggregated plots (figure from Elhamifar & Sastry, 2015)	13
Figure 5: Feed forward ANN (Xu & Milanovic, 2015).....	19
Figure 6: Mockup showing disaggregation workflow	23
Figure 7: Proposed Disaggregation workflow with external datasets, notifications, and quorum	27
Figure 8: On vs. off durations for common home appliances (figure from Hart, 1992)	30

Lists of symbols, abbreviations or other

- AC: Alternating Current
- ANN: Artificial Neural Network
- DC: Direct Current
- DTW: Dynamic Time Warping
- HMM: Hidden Markov Model
- MFCC: Mel Frequency Cepstral Coefficients
- NILM: Non-Intrusive Load Monitoring
- RFID: Radio Frequency Identification
- SVM: Support Vector Machine
- TTL: Time To Live

1. Introduction

Property ownership comes with many associated costs. Some costs associated with homeownership are easy to budget for due to low variances in amount and frequency. Property taxes would be an example of a budgetable expense. Maintenance is at the other end of the spectrum because it is sporadic and varies according to several factors such as property age, usage, and weather patterns (Kolokotsa et al, 2011). Utilities hold the middle ground of expected frequency and varying cost. Depending on occupancy and appliance usage, utility costs may vary greatly on a month-to-month basis.

Of the top costs associated with homeownership, utilities rank fourth after mortgage payments, insurance, and taxes (USAA, 2013). Although electrical costs comprise only 4% of a Canadian's average daily budget, trends show this has been increasing annually with rising energy costs (StatsCan, 2011). North American homes may see similar pricing to European countries if our electrical consumption issues are not addressed. The average household in Denmark spends 4 times more on electricity compared to Canada (IEA, 2012). Canadian electrical prices increased 9% overall from 2010 to 2014 (Hydro Quebec, 2014). Trends show utility costs will continue to rise.

Rising energy prices and increasing energy demands have created a trend of energy conservation initiatives. New products, such as Belkin's Conserve Socket (Belkin, n.d.) are being created to autonomously turn off power to devices that are not in use. There are also a slew of sensors to control lighting, HVAC and blinds (SmartHome, n.d.). The common goal is to reduce energy consumption. The autonomous category of devices is paralleled with habit-forming devices that provide information in order to persuade consumers to conserve.

Forming conservation habits using real-time visual feedback at the appliance can influence consumers to save 12% on energy costs annually (Armel et al, 2013). When real-time feedback is paired with incentives, these savings can be increased to 32% (Peterson, 2007). Smart metres are an example of devices that communicate real-time power consumption through the display of dollars per hour or

kilowatt-hours. The smart part of the metre is a misnomer to consumers because they don't directly enable users to reveal devices that consume energy. This is due to configuration of these devices relative to the mainlines, which only allows for aggregate readings. Users can only receive an overview of consumption for the entire premises.

<i>Table 1: Comparison of disaggregation benefits by group</i>	
Perspective	Benefits
Consumer	Conservation, pricing incentives, financial savings, appliance replacement, automation, property valuation, usage tracking, feedback on initiatives, auditing bills
Utility	Pricing considerations, remote metre readings, regional load tracking, conservation recommendations, distribution, planning
Manufacturers	Appliance innovation, lifespan assessment, fault detection, warranty assurance, efficiency validation, gathering customer knowledge
Building Design	LEED certification, energy efficiency assessments throughout the year, automation, research data
Government	Energy policy decision making, census studies, forecasting, funding, incentives

Data made available by monitoring an electrical line allows for further analysis of what devices are causing consumption. The alternating current power system uses a frequency much like audio or light. For sound to be perceived, fluctuations must occur within a frequency ranges that can be perceived by the human ear. Creating software and hardware to track a speaker in a room with many people is a task comparable to disaggregation. If a single person is talking, it's a seemingly simple task. When multiple people are talking at the same time, the same task begins to compound in difficulty with every additional concurrent speaker. The issue associated with tracking an individual speaker in a room where multiple people are talking is known as the *cocktail party problem*. Furthermore, disaggregating multiple potential sources of speech is known as *blind source separation* (Kyan, 2015).

The method of deciphering an individual speaker is comparable to tracking electrical usage of a single device from an aggregate reading. Electricity is used by devices at varying magnitudes, much like varying volume when talking. For the electrical spectrum, we are able to use a comparable set of inputs. Vocal tone, accent, modulation, and sentence structure all contribute to a vocal signature (Bronkhorst, 2000). For the signature of an electrical appliance, these contributions would take the form of harmonics, current draw, duration of usage, and transient states. Measurement quality, accounting for multiple characteristics, and classification algorithms are components used to improve recognition accuracy.

Context can also be used as a method for disaggregation. For listeners, context is provided through lip movement, gestures, and eye direction. Being able to see the speaker can help simplify the isolation of their dialogue in a noisy environment, much like knowing a user's location simplifies the probability devices within their proximity are in use. Date, time, weather patterns, number of occupants, location of occupants, and recurring habits can all be used to provide adjustments to probabilities. Context can enable educated guessing. Not only does it help solve decisions when choosing several highly probable options, it can be used to improve processing efficiency by reducing the number of options to account for.

2. Approaches to Electrical Monitoring

Initial readings are obtained through paper statements, power plugin metres, smart metres, or induction clamps. Each method has its pros and cons. Induction clamps (as seen in figure 3) are a non-intrusive means to monitoring with low costs and high sampling rates. Smart metres are gaining adoption rapidly because utility provides benefit from remote monitoring (Navigant, 2014). Plugin metres, such as a Kill-A-Watt metre (P3 International, n.d.), are highly accurate, but unattractive due to their cost and manual tracking requirements. Bills are gateways to other monitoring methods since they alert consumers to the growing costs of electricity.

Table 2: Comparison of tracking methods

	Accuracy	Sample Rate	Installation	Adoption	Cost
Bills	0 devices	Billing Cycle	None	Ubiquitous	\$0
Plugin metres	One per plug	~1s	Easy	Low	\$20-\$60/ea
Smart metres	Large appliances	1s+	None	Very High	\$0
CT Clamps	Potential for all	>1 Hz	Hard	Low	\$150+

2.1 Bills & Statements

The general low-tech method to monitor electric consumption is through the review of a periodic bill. The minimum content a paper statement conveys is kilowatt usage and associated charges. Newer statements may provide daily trends and graphs that communicate usage over the billing period. Considering the delays involved with receiving a paper statement, communicating unusually high consumption will often be after the fact. Bill recipients are unable to be proactive about their usage and therefore play a reactive role in conservation. Therefore, bills play an initiating part in influencing conservation.

Once electrical expenses are realized, conservation efforts require further investigation. Property owners first spend time reviewing what appliances are onsite and then begin reviewing historic usage if available. Without some form of logging, it becomes necessary to employ a proactive monitoring strategy to review which devices are in use and compare observations to resulting bills. The time requirements to obtain feedback and inaccuracy make for an inefficient approach to monitoring.

2.2 Plugin Metres

At the device level, power plugin metres provide the most accurate measure of real time consumption. A typical plugin metre ranges from \$30-\$80 each and are widely available. Typical

statistics provided by these metres include forecasting, cumulative kilowatt hours, volts, amps, watts, hertz, volt-amps. Many libraries in North America offer the Kill-A-Watt plugin metre for loan. With an accuracy of 0.2%, it is the most accurate method of logging device consumption (P3, 2015).

Vampire power or parasitic load occurs when a device is off, yet still plugged in and drawing power (Kim et al, 2011). Although the amount of power consumed is relatively low, having many devices drawing low amounts of energy can add up to a significant load. Plugin metres have the added benefit of sensitivity since they can detect standby usage.

The primary downside of a plugin metre is cost. With the number of devices in a typical household or business environment, it would cost thousands of dollars to install enough plugin metres to monitor each (unique) device. Should a building manager decide to accept the costs involved, the monitoring and logging is typically a manual task with these devices. A plugin metre excels in creating an accurate real-time snapshot of a particular device. Typical plugin metres do not provide statistics other than a kilowatt average and real-time load information. Therefore, they are not an effective long-term solution to conserving energy.

Assuming only a handful of plugin monitors were implemented, they would typically be used to monitor major appliances. Users would typically miss out on the many low energy devices. In some cases, the consumption of these devices shall outweigh the few energy hungry devices in terms of overall energy consumption. This distribution is known as the long tail (Anderson, 2008). If aggregate data is being tracked, it is possible to determine the scope of this amount by simply subtracting the monitored devices from the total aggregate. However, this does not account for any unknown devices that are not anticipated as heavy consumers.

Monitoring issues aside, another downside of plugin monitors is the additional energy these require to operate. Tracking for purposes of conservation by these devices will increase consumption slightly. The use of many plugs at a location would increase consumption and may negate the value of

their use. Plugin metres do consume some power to operate. If many are left plugged in for an entire year, the total consumption to run the metres needs to be considered.

Considering the equipment cost, consumption, and labour associated with tracking plugin monitors, it does not make sense to use them to track the majority of devices in a building. Similar devices such as power bars with built in digital monitors offer the same drawbacks. The benefits in terms of precision are not important enough to justify their usage.

2.3. Smart Metres

Smart metre adoption has been growing steadily at about 5% per annum in North America and projected to reach 80% penetration by 2023 (Navigant, 2013). Utility providers save money because smart metres enable them to remotely receive metre readings. The traditional method involved an employee physically visiting a location to manually log metre readings. Readings would then be input into a database for billing purposes. Due to the labour and time involved with the traditional process, smart metres not only help the bottom line, they also offer a better resolution to readings so utility providers can understand usage by region. This information is used for planning, marketing, and forecasting.




Figure 1: Smart metre penetration by region

Utilities have been using smart metres to overcome the costs and lead time associated with technicians physically visiting locations to track analogue metres. There is enough incentive to push this technology to their customers. Therefore, there is no impact to their customers in terms of cost, network infrastructure, or additional hardware within the premises.

Installation of a smart metre takes place where analogue metres are installed. In commercial and industrial buildings, there may be several locations where metres are installed. These additional metres are at entry points of mainlines and sometimes sub metres that are installed to gauge energy use by area or industrial equipment. A utility is concerned with the mainline metres for billing purposes, not submetres.

There are many variations to smart metres, but the most prevalent ones are retrofitted onto existing analogue electrical metres. These adaptors convert analogue signals into digital signals for dispatch to the utilities servers for billing and tracking purposes. The data being transmitted by a smart metre represents the aggregate electrical consumption of a building. Since aggregate data is being made

available by these devices, it appears ideal for the data to be analyzed for disaggregation. There are however limitations to the data being made available due to sampling resolution.

Smart metres produce data at intervals ranging from 1 hour to 1 second. The intervals are dependant on manufacturer configurations and are limited by hardware and firmware (Armel et al, 2012). Lowering the interval is primarily to reduce network loads and hardware costs, but also has a relation to privacy. Higher frequencies would reveal customer habits, which has been cited as a concern by privacy groups advocating against smart metre technologies. Privacy issues are covered extensively in McLaughlin et al's (2011) paper on *protecting consumer policy from electric load monitoring* and Lisovich & Wicker (2008). Privacy concerns are not the focus of this paper.

In terms of network load, the infrastructure currently used to transmit data has the bandwidth for increased intervals over 2 kHz range (Armel et al, 2012). The restrictions on this are due to hardware at the analogue to digital conversion, processor, memory, communications interface, and home area network. Within 1 Hz to 2 kHz, the restrictions are primarily due to firmware, not hardware, because compressing data batches alleviates the need to strain communications interfaces. Surpassing the 2 kHz sample rate allows for medium order harmonic devices to be detected, including toasters, computers, and televisions.

Due to the low frequency of data intervals (greater than 1 second), the decision of which algorithm to use comes with limitations regarding the type and number of devices that can be detected. Smaller devices in use are left out while only major appliances are detectable. Very few devices can be detected with a high accuracy causing a major limitation with this technique (Armel et al, 2010).

The primary benefit of pursuing smart metres for disaggregation is their prevalence in developed countries. Utilities are the primary beneficiary of the data these devices produce. Since the cost and installation are handled by the utility, it places no burden on the customer aside from a brief power outage during installation. Smart metres are already placed on the main lines entering a premises, so it would be

ideal to modify the transmission interval for disaggregation. As an added benefit and incentive to customers, utility providers can produce relevant statistics that can influence conservation initiatives from their customers.

Daily pricing schedules are being replaced by peak pricing methods. Utilities are shifting towards variable pricing strategies to initiate customer driven conservation. In order to communicate changes to pricing, short-term notifications through email and mobile applications are paired with real-time pricing indicators on a smart metre (Peaksaver, n.d.). These methods are enabled because the smart metre allows for customers to be more informed through digital technologies.

2.4. Induction Clamp

Also known as current clamps, clamp metres, or current probes. Induction clamps are an ideal sensor for non-intrusive load monitoring because they support a high sample rates and do not require mainlines to be intrusively modified in anyway. These devices are comprised of a split ring that can be easily opened by hand and clamped around a cable. Once installed, data generated from the clamp is transmitted to a measuring transmitting unit (MTU). The MTU can then transmit readings through a serial interface, powerline Ethernet, Ethernet, or wireless transmission.

An AC clamp metre generates data using a current to voltage converter, also known as a current transformer (CT). The principle of variable flux is used to generate signals that in turn generate data (Silva, 2005). Fluctuations in current can be thought of as a relative motion to the point of observation by the clamp. This rapid electrical movement creates a magnetic field around the cable. The clamp's metallic core is position around the cable and therefore is able to observe the alternative current of electricity. CT signals to an analogue to digital converter for logging and display purposes.

In North America, homes usually have 2 main lines entering a building. These lines distribute electricity 120V of electricity and are 180° out-of-phase from one another (Hart, 1992). Appliances consuming 120V are connected to breakers that alternate between the 2 lines in a breaker box. 240V

appliances are wired in an unbalanced fashion resulting in varying consumption on both lines. Therefore, a current clamp is needed to monitor each line and used additively to provide an accurate aggregate measurement.




Figure 2: Breaker box with 2 phase configuration



Figure 3: Induction clamp installation

Induction clamps are ideal disaggregation monitors because of their price relative to other options. They also come with the benefits of easy installation. One of the most important aspects to consider is a clamps ability to provide a sample rate of greater than 1Hz. Compared to smart metres, clamps don't require any hardware modifications. The importance of resolution is discussed in the section 4 of this paper.

3. Opportunities with Disaggregation

Having detailed data on energy consumption on bills account for a 10% reduction in energy usage (Darby, 2006). Direct feedback through a real-time monitor can reduce consumption by up to 20% (Faruqi et al, 2010). Although having a metre provide live consumption readings can help reduce energy consumption, it does not reveal the specific devices that are contributing to the total reading. Without device level information, conservation efforts cannot be targeted unless further investigation is conducted. Combining real-time monitoring with specific device information and automation enables savings ranging from 15 to 40% (Mercier & Moorfield, 2011). This value does not take into account the time saved by not having to investigate aggregate consumption readings.

Using disaggregation information to notify users that a device has been left powered on would be another major benefit of disaggregation. Through a mobile notification platform, messages could be received for devices that are being used outside of their standard usage duration or time of day. A projector or lights left on overnight could be dealt with in sooner, saving on energy costs and maintenance. Combining smart, network enabled devices with remote notifications can allow for remote shutoff, especially in the case of computers.

If devices have been turned off in a certain area, it may indicate users are not present. This information could be used to turn off air conditioning, heating, and lighting after a grace period. Remote shutoff can be enabled through networked devices or smart electrical terminals. The number of network-enabled devices is growing at a compound annual rate of 25% (James, 2014). The emersion of devices that support network communications will add to the cache of controllable devices.

The opposite of remote shutoff also holds true. Devices could be linked together in workflows to enable group activations. Such activations could be performed using device states by interpreting disaggregated signatures. This would allow automation that is driven by device processes rather than on/off states. For commercial applications, there may be an opportunity to create interoperable systems.

For example, a printing press changes speed once a print run is started. The initial make-ready process is much slower because it is used to prime inks and register offset images. The difference in speeds result in a change in power draw and could be used to run converters or have logs created in job tracking systems.

Instead of having many devices on standby, disaggregation could allow a single device to trigger start-up at the level of an electrical circuit. A single circuit often distributes power to a zone or room in a building. Zone-based control would be possible, allowing entire areas to activate their respective devices. For example, a swipe card lock could enable current to all devices within a room. Networked devices could be used to allow complex logical states that are workflow-based rather than zone-based.

Certification programs, such as Leadership in Energy and Environmental Design (LEED), assess buildings on their energy efficiency. Measurement and validation are used as rating criteria. Active monitoring and disaggregation allow for both aspects to be realized. Certification programs are only validating a premises for a snapshot in time (Lee & Burnett, 2008). Maintaining an electrical measurement and validation system can be provided using non-intrusive load monitoring (NILM). Through disaggregation, a building's occupants and devices could also be assessed to dictate conservation initiatives. For example, without device specific data, the specific benefits of upgrading or replacing equipment is unknown.

Through the analysis of device specific energy signatures, the devices state can be revealed. Not only the operational state (on/off), but also the device's efficiency over time. Poorly operating devices consume more electricity due to friction and degrading components (Xudong et al, 1998). Increased consumption at various cycles could be used to indicate the specific components that are failing before equipment is permanently damaged.

4. The Building Blocks of Disaggregation

Disaggregation is made possible due to the unique electrical consumption patterns of devices. A consumption pattern can be based on *power draw*, *functional or transient states*, and *frequency of use*.

Each of these characteristics contribute to a unique device fingerprint that can be used by algorithms to isolate the device from the aggregate consumption reading, which is often captured through the use of induction clamps on a building's main electrical lines.

For example, a typical washing machine has a wash, rinse, and spin cycle. Each one of these cycles have a discernable consumption pattern. These changes in consumption reveal changes in a device's operation mode and are referred to as transient states. Algorithms can use transient states to determine the current device operation mode and instigate the past and upcoming cycles. There will be some consumption variances between manufacturers and how a machine handles common programs, such as; delicates, permanent press, cottons, and so on. Variations would require training for a database to capture the appropriate device signatures to dictate disaggregation accuracy.

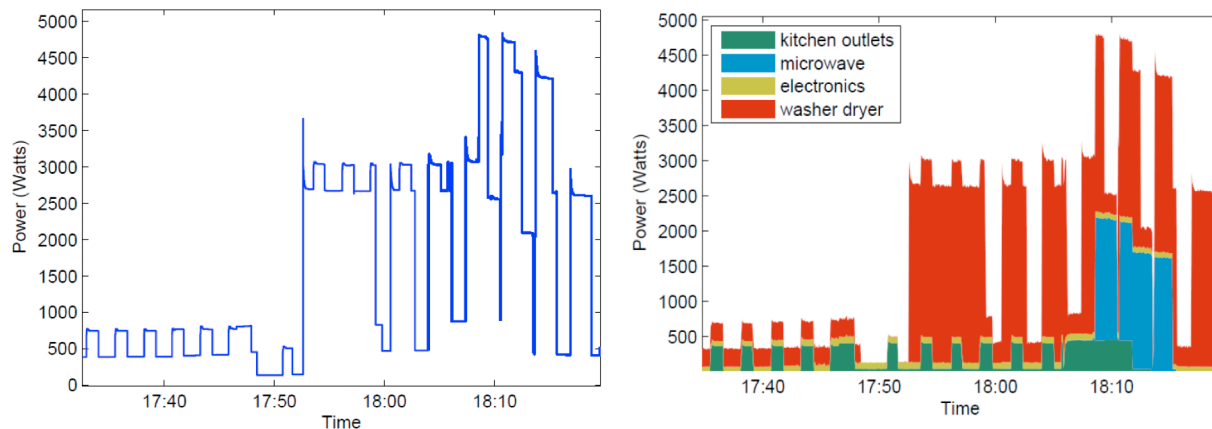


Figure 4: Aggregate and disaggregated plots (figure from Elhamifar & Sastry, 2015)

A good way to visualize disaggregation is through the use of blocks. Each block represents an individual device's signature. Blocks come in various shapes, sizes, and colours. The shape of a block would relate to the consumption as it varies over time. The shape represents the magnitude of consumption. Colours can be assigned to show device type. Aggregate consumption can be seen as a wall containing many coloured blocks. There is often more than one configuration that can be used to build the

wall. Disaggregation algorithms try to build an accurate representation of the wall when only being provided a silhouette of the wall.

Whenever a new device transitions in terms of its consumption, a peak or trough occurs. These changes relate to the device's components turning on or off. It can also relate to a change in functional states. For example, a washing machine switching from a wash to spin cycle. Changes in consumption are typically drastic for appliances with large power draws. Edges are formed by the sudden functional changes and create an ideal starting point to find matching signatures.

The total initial draw of a device will have an equal reduction when usage ends. If a 60W bulb is turned on, the aggregate usage will reveal an increase of 0.5A on a 120V circuit. This consumption will be sustained for the entire duration the bulb is left on. Once turned off, there will be a 0.5A reduction in the aggregate consumption. If the bulb is the only item in use, it's very easy to see the point in time where it switches on and off.

When multiple bulbs are turned on or off, there is no way of discerning which exact bulb was involved. The complexity of disaggregation is compounded when many 60W devices are in use or there are appliances that have 0.5A fluctuations in their consumption as they change states. Unless additional information is provided to a disaggregation algorithm, accuracy will remain low for devices with similar consumption signatures.

Hart (1992) proposed the use of a Zero Loop-Sum Constraint (ZLSC). The premise of ZLSC is that all devices will always have an equal increase and decrease in power consumption throughout their use. If a device draws 10W when in use, then we will observe a -10W drop when not in use. This is applicable to all transient states a device may have. Comparable to the conservation of energy law, ZLSC can be used as a method to account for multiple devices when disaggregating.

4.1. Power Draw

Power draw is based on the appliance's consumption. An algorithm may use additive consumption to detect a device for the duration it is on. This works for devices that have a consistent power draw for each usage case.

Another method is to use the spike in consumption that occurs when a device turns on. This is often found with devices that require more energy to overcome inertial forces such as motors and heating elements (Hart, 1992). These types of devices can be modeled as linear. Real and reactive power can also be used to detect devices.

Real and reactive power are two distinguishing characteristics of a device. Real power occurs when resistance dissipates power while voltage and current remain in phase (Kelly & Knottenbelt, 2011). Reactive power in contrast borrows and returns power to the AC source but the voltage and current are out of phase from each other.

4.1.1. Real Power

Real power is found in devices that are resistive or use electricity with a resistant element to produce heat or light. When such a conversion occurs, the characteristics of this conversion can be detected and used to model to a device accordingly. Sub-characteristics of a device's real power can be noted over time. These can assist when differentiating devices from one another. For example, resistance can increase or decrease over time. A light filament's resistance increases in relation to heat while a heating element stabilizes once the threshold temperature is reached (Hart, 1992).

4.1.2. Reactive Power

Devices that step down or drop voltage exhibit reactive power. The unit of measurement is Volt-Amps-Reactive (VAR) and is represented by the letter Q. Reactive power occurs in inductive loads that are stored and discharged. Devices containing motors, transformers, or solenoids exhibit reactive power

(Piotrowski, n.d.). Although the energy used is stored and dispelled back into the system, it causes returning electricity to be out of phase.

4.1.3. Harmonics

Certain devices inject electrical anomalies into the alternating current power system known as **harmonics**. Electricity in North America is delivered at 60Hz and harmonics occur at intervals that are multiples of 60 (Ling & Eldridge, 1994). There is some distortion inherent to electricity that is created during generation, but this varies between 1% and 2% (Sankaran, 1999). Harmonics are created by devices such as battery chargers, computers, and motors with variable frequency drives. These devices are described as non-linear because the voltage distortions created are at particular frequencies above the 60-hertz range. Because the distortions occur at frequencies that are multiple of the standard, they can be more readily isolated and used to characterize a device.

A Fast Fourier Transform (FFT) allows for information in the time domain to be adapted to the frequency domain. Through this method, we can reveal devices that produce harmonics (or any other characteristic frequency signature). Electrical data being captured must be at a sample rate much greater than 60Hz for this method to be useful. Due to the sample rate being higher than what is generally available through monitoring devices, it is difficult to rely on harmonics for signatures unless specialized monitoring hardware is used.

4.2. States

A *state* (in the context of energy disaggregation) can be described as a steady consumption pattern over time. A standard light would exemplify a binary state (dual state of ‘on’ or ‘off’). Multiple states can also be found in devices such as room fans, where a user is able to select different speeds. As long as the fan’s motor functions in a steady state, a pattern can easily be deciphered from aggregate consumption data. Transient states are found when a device has multiple cycles. Washing machines and dish washing machines both have multiple states in each of their programs. The difficulty in

disaggregating devices with transient states is due to the length of each state being dynamic in terms of duration and power draw. To overcome accounting for transient states, ZLSC can be used to balance the sum of power changes by ensuring total of all state transitions are equal zero.

4.3. Frequency of Use

Frequency based identifiers work well for devices that contain a switched-mode power supply. These are byproducts of modern devices that release harmonic distortions into the electrical grid within a building. They may also have transient states and programmed cycles. Examples of these devices are computers and low voltage DC devices that convert alternating current. In order to analyze the consumption patterns of switched-mode power supplies, the time-based consumption readings must be converted into frequency-based readings. Using a Fourier transform and a sufficient sample rate to reveal harmonics are requirements for this to be possible. Ideal rates are 60Hz for low-order harmonics and over 1MHz for high-order harmonics (Armel et al, 2012).

5. Algorithms for Disaggregation

5.1. Factorial Hidden Markov Model (FHMM)

Developed in 1906 by Andrey Andreyevich Markov, the Markov Model is used to reveal patterns within stochastic or chaotic data chains. The key characteristic of a data chain is its capture over a period of time. The preceding and succeeding data fragments have a relationship and therefore can be used to create a baseline for comparison. It is an opportunity to filter outliers and develop the underlying pattern that is trying to be modelled.

A *Hidden* Markov Model (HMM) is used for calculating multidimensional probabilities when there are a number of observable states (Kouemou, 2011). Each state is a part of a chain of segments. We are analyzing an excerpt of the overall data chain and comparing data segments to future and past segments to extrapolate states. Therefore, we have no baseline function from which to build a start and

end state. Instead, we are modelling a function that comes as close as possible to the real function. The calculations are based on probability and therefore are meant to infer what is observed from the data collected. The Viterbi algorithm is then used to predict the outcome or resulting state (Ghahramani & Jordan, 1997).

An analogy for the HMM would be to consider a taxi picking and dropping off passengers at various locations throughout a city. The pickup and drop-off locations are known along with the time taken to complete the journey. The possible route taken for each journey is unknown but can be determined by calculating the time for each possible route. Approximate route time for each variation is then used to calculate probabilities for each possible route according to how close it is to the expected travel time.

HMM is an efficient yet simple machine learning method for classification of a single appliance. Using HMM is unrealistic because multiple devices need to be found in an aggregate reading. HMM works well for single device searches in the aggregate. For energy disaggregation, the *Factorial Hidden Markov Model* (FHMM) is employed because it can account for multiple independent relations. FHMM avoids the complexities involved with using HMM to reveal several underlying signatures. For FHMM, each signature or device is considered to be a chain in parallel to one-another at the cost of precision (Mackonin, 2009). Hidden state estimation is performed using a process such as Gibbs sampling (Kim et al, 2011). The resulting accuracy for this method alone is approximately 65% (Kolter et al, 2011). To improve disaggregation accuracy, it is not practical to rely solely on FHMM. A characteristic of FHMM is that a state occupancy duration is constrained to be geometrically distributed (Kim et al, 2011).

Between 69-98% accuracy on 10 homes was achieved using FHMM, CFHMM, CFSHMM, EM, and MLE (Kim et al, 2010). Using variations of FHMM, DFHMM, AFHMM, and a custom Additive Fractional Approximate algorithm, 71% accuracy can be achieved (Kolter et al, 2012). Another combination of FHMM, FSM, and VAST could be used to achieve an accuracy of over 87.2% on major appliances (Kolter et al, 2012).

5.2. Artificial Neural Networks (ANN)

Neurons are specialized nerve cells found in the brain that use electrical and chemical signals. The way neurons are connected to one another has inspired the machine learning method known as ANN. Modelling a neural network requires a real value weight, an activation function, input connections, one or more hidden layers, and an output layer (Xu & Milanovic, 2015). The real value weight is a number between 0 and 1 inclusively and is used to quantify connections between neurons. Activation functions allow neurons to be triggers if the sum of input connections meet a predetermined threshold. The output from a neuron is used as inputs to neurons connected to subsequent layers. The number of layers is determined through training or manually adjusted. The output layer is the final layer in the neurons where an observable output is provided.

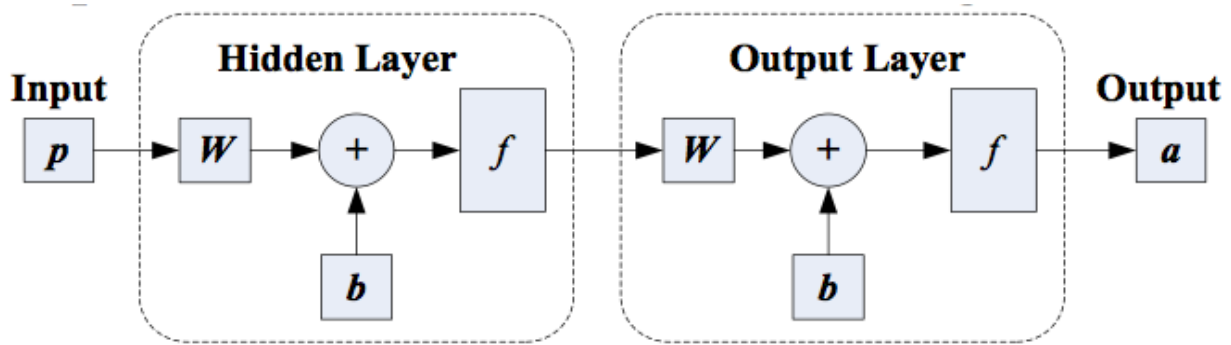


Figure 5: Feed forward ANN (Xu & Milanovic, 2015)

After outputs are observed, training must take place to improve ANN accuracy. The backpropagation algorithm is used to train ANN (Siau, 2004). This measures the expected output relative to the resulting output. Errors at each neuron are measured and provided an adjustment factor. Training is a manual procedure that is required to correct errors. As devices are added to the system, further training is required. The result is disaggregation accuracy that greatly drops below 40% when greater than three devices are added (Chang et al, 2009).

5.3. Support Vector Machines (SVM)

SVM is the strategy of determining the line with the greatest margin between recorded signatures. This is known as optimal line separation. By determining margins, signatures can be differentiated during training. For actual data, margins are calculated again to classify signatures found in aggregate. SVM performs this task efficiently for small datasets (Marsland, 2009).

Figueiredo et al (2012) was able to achieve 96.4% accuracy using SVM for 6 household devices using a low sample rate. By testing various SVM and k-NN combinations, they were able to further improve the algorithm to 99.1% using a combination of 5 nearest neighbors (5-NN). Due to the low sample rate, the algorithm worked well for major home appliances, but further testing would need to be done to assess their algorithm on faster sample rates and a larger range of devices.

5.4. Nearest Neighbour (k-NN)

Through the measurement of distance between data, the nearest neighbour algorithm rates data to training data. Therefore, k-NN is an algorithm requiring supervision to be implemented. When classifying data, k-NN uses a distance calculation to vote which signature is the closest match. The k-factor is used to determine the search output. If there are several signatures that are a close match, a k-factor of 3 would return the single closest signature. When set to 5, the result is based on the majority vote of the closest neighbours (Makonin, 2012). On its own, k-NN does not perform well for disaggregation. It can detect major appliances, but requires extensive training and adjustments to k.

5.5. Dynamic Time Warping (DTW)

Developed for linear sequences, DTW originated as a speech recognition algorithm (Sakoe & Chiba, 1978). DTW approaches comparison through the measurement of similarity between sequences. Warping refers to the non-linear comparison that is made possible by the algorithm. The shape of a series in terms of magnitude (y-axis) is what matters more than the location of a point in time (x-axis). This

makes DTW ideal for speech and handwriting recognition, and signal disaggregation. Implementing DTW for disaggregation requires the algorithm to run after event detection is completed. A recognition accuracy of 85% can be obtained using DTW for disaggregation (Elafoudi et al, 2014).

6. Disaggregation Workflow

6.1. Measure Aggregate

Measurements are obtained at a defined interval in aggregate from the mains. The sample rate is dictated by the hardware and software in use. Data is often stored locally and polled or posted to an external server in batches for processing (TED, 2010).

6.2. Store Aggregate Reading

Although raw data may be captured at a high frequency, it may be unnecessary to store high-resolution data once harmonics are revealed. The primary reasons for this are overhead and the user's needs. Storing data in 1-second increments is sufficient for review and targeting large appliances. The actual storage procedure should batch intervals together, compress, and transmit data to an external database to reduce network traffic. This is especially important when the database is being stored offsite. Therefore, a local buffer is needed to batch intervals together and perform the initial pass.

6.3. Initial Pass

Raw aggregate readings require several operations before being run through segmentation. The first is normalization, which is used to compensate for problems that occur during the delivery of electricity. These include external factors such as fluctuations in voltage. Another aspect of normalization is splitting the two legs of the 120V circuit that enter as 240V in North America. The second step of the initial pass is to remove noise and amplify the readings so the segmentation process can be more accurate.

6.4. Data Segmentation

Edge or event detection is used to determine when devices on the aggregate readings are changing states. This information is used to break signatures into pieces that are comparable to the existing signatures in the database. It is beneficial to store segment timestamps for recall when a device signature spans multiple segments due to transitions in states.

6.5. Signature Matching

Segments are processed through cluster analysis (Kouemou, 2011). In this stage, FHMM or comparable algorithms are used to disaggregate individual device states. Hints from the database of existing signatures can be used to detect transient states and harmonics. Devices should be matched within a tolerance of $\pm 10\%$ because of the expected voltage variances. This may be adjusted depending on how well the initial pass removes variances in advance.

Once a match occurs, the device information is stored into the database. This information is used to educate the user when they are reviewing their consumption statistics.




Figure 6: Mockup showing disaggregation workflow

7. Furthering the Workflow

Additional steps can be taken to improve signature classification and conservation efforts. The goal is to grow the database of signatures over time in order to create a system that is increasingly adaptive and accurate. Such an improvement would allow for an opportunity to educate the user in real-time on their behaviours and develop strategies that guide conservation efforts according to detected trends.

The standard protocol is to turn all devices off and classify the signature of each device, or use a separate piece of hardware, such as a plug, inline with the device's circuit. These standard methods are time consuming and require a user to be actively involved in the process of capturing device signatures. It is important to also consider that the property would experience downtime of all equipment while training is taking place. Any additional users would need to coordinate their use of electricity so as to not contaminate the trained signatures.

This paper proposes an asynchronous, less involved approach to classifying unknown signatures. Some training is necessary especially for the signatures of major appliances that are known to draw significant electricity. Fridges, stoves, washing machines, clothes dryers, microwaves, and air conditioners are examples of these appliances. These major appliances account for 60% of residential energy consumption (Hart, 1992). Providing these signatures have been captured, we are left with the remaining signatures of several low-power devices and a handful of mid-power devices. Low-power devices consume less than 50W, such as: lighting, security cameras, proximity or IR sensors, and key fobs. Televisions, sound systems, game consoles, and projectors constitute mid-power devices which consume less than 500W.

The storage of unmatched edges can follow the storage of matched signatures. Since each edge describes a device transition or on/off state, it can be used as an opportunity to classify a device for future comparisons. A user would need to intervene and provide device information to the database. This step can take place real-time when a new device comes online or when a user is reviewing their aggregate usage over a period. All newly categorized signatures are centralized for other users to benefit from the knowledge contributed by the community.

A crowd-sourced approach to signature classification would require the following considerations:

1. Calibrated hardware
2. NILM Hardware with greater than 1Hz resolution

3. Details on device characteristics
4. Consensus

Calibration is an important step in configuring the device for use in a crowd-sourced approach. The expectation is that there will be a 10% variation in the voltage being delivered to the premises (Hart, 1992). If hardware is incorrectly calibrated, the variation may be amplified. Because edge detection requires a threshold to reveal a device changes state, variations in delivery may also set off false positives in the edge detection algorithm resulting in invalid signature captures. It is not only inefficient to try and classify an invalid signature, it is also detrimental to the entire disaggregation system because an increased number of signatures requires a greater amount of processing to search for a best match.

By having the user provide *device characteristics*, the classification algorithm can use this information to predict on/off and transient states. The features can also be used to infer how often the device is in use, usage duration, standby state, autonomous usage over a period (turns itself on for maintenance), and the number of states. The makeup of a device would reveal parts that contribute to the consumption pattern, such as motors or resistive components that give off harmonics. A combination of device harmonics, usage durations, power draw, and operating states allow an algorithm to develop a device profile that can be used for search, comparison and classification.

Consensus or quorum allows a signature to be vetted before being used as a truth for all users. Some users may improperly classify a signature to a particular device. When a signature is initially classified, it is saved as a hint to direct other users in the validation process. Should other users agree with the classification, the probability of the signature being an accurate match increases. As more data is collected, the types of matches found would increase, but may also cause multiple matches for a single unknown signature. Therefore, the number of users to reach quorum to classify a device is an important consideration. The opinion of 4 users can be taken to emulate a single expert user (Snow et al, 2008). Therefore, when 4 confirmations of a signature are received, we can use the match to automatically classify the device for other users encountering the same signature.

There will be cases where consensus improperly matches a signature. Users coming across this issue will need to reclassify the device. The change would result in another hint if the new classification does not exist in the database. Once again, 4 users in total will need to agree with the classification before the device will be automatically classified for future users. The variations between the 2 similar signatures will need to be amplified to improve categorization in the future. In a case where the consumption patterns do not vary for the devices, the system would need to recall past classifications (if any) from the user and use those as a basis for future categorizations. This means if 2 devices (A and B) have the same signature and a user classifies their signature as device B, all future matches would automatically be classified as device B without requiring further intervention by the user.

The outcome of this process is to provide the user with historic trends for the device, allowing them to intervene and conserve energy. Having accurate classification when there are multiple variations of the same signature reduces issues associated with having to search and locate devices in a property. In turn, users are enabled in taking action because the information needed to make conservation decisions is clearly communicated.

7.1. Workflow Using Crowd-sourcing

1. Unknown signature detected (or possibility thereof - due to edge occurrence)
2. Signature stored for classification as a hint with timestamp
3. Device details are input by user
4. Quorum is achieved
5. Autonomous classification occurs for all users
6. New hint is created when user tags a wrongful classification
7. Quorum for new hint is achieved
8. Variation in 2 potential matches is determined and stored
9. Autonomous classification occurs for all users

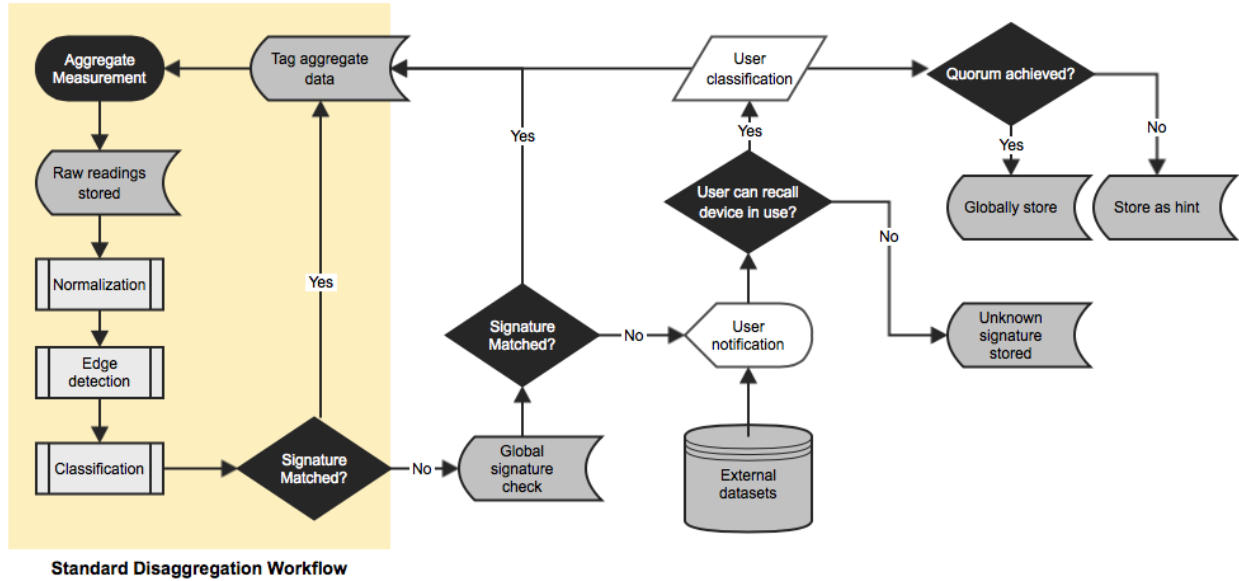


Figure 7: Proposed Disaggregation workflow with external datasets, notifications, and quorum

Although the user may not know the particular consumption pattern of a particular device, some characteristics can be provided. The power supply often denotes the voltage and amperage. The user could also classify whether or not the device has multiple modes of operation, number of usage states, resistive, or releases harmonics. To simplify classification, the device could be categorized so these properties can be inferred. The user could also determine their typical usage in regards to time-of-day and duration. A brief survey of the components (motorized, resistive, computerized) can lead the database to capturing the device in the future. Without some form of training to capture a device signature, a survey on its own will not result in accurate disaggregation. Instead, survey results can be used when uncategorized signatures are detected.

Users contributing categorization information are able to derive context through the use of timestamps. The moment in time where a device is in use provides users with an opportunity to recall what was in use. This coupled with disaggregation information of matched devices can increase the chances of accurately matching the device in use. Without having user context present it is much harder to determine what the device is purely based on an energy signature. Therefore, the relationship of time and

device usage allows the user to provide their understanding of their environment in order to categorize an unknown device.

Another opportunity to further device detection is notifications through a smartphone or other network enabled device. Asking the user as soon as an unknown device is detected would improve categorization while adding signatures to a centralized database. Other users can benefit from that shared intelligence. Finally, an interface could be created to allow the user to categorize unknown devices during a review process.

8. Integrating External Datasets

To improve the probability of categorizing a detected signature, contextual information from other sources can potentially aid in improving accuracy. In order to provide context to an algorithm, an external dataset provides hints that can be used to change the weights of probabilities associated with particular signatures (Kim et al, 2011). The simplest dataset could rely on the probability a particular device may be in use at a given time of day. This is an example of what George Hart describes in his algorithm that references time of day, duration, and idle time (Hart, 1992). There are many other datasets that can be cross-referenced to improve disaggregation. This paper will analyze the process of creating datasets and concepts of how these datasets can be applied to improve accuracy. Additionally, the process of solving unknowns arising from the disaggregation process will be discussed.

8.1. User Habits and Trends

Recurring events occur at all homes and businesses throughout the day. These trends are more apparent for homes and businesses with set work routines. Tracking trends and cross-referencing them to electrical consumption would be one method to create an external dataset. Any datasets that relate influences on user behaviour to consumption can be used to infer device usage. Certain datasets are directly related to user actions while others are linked to actions. For example, a direct result of turning on a smart TV would be a connection to the network to check for a software update. The action of turning on

the TV could therefore be cross-referenced with network traffic and validated to the aggregate reading. Although this would verify the TV is in use, the TV does not always make this request. Therefore, a separate dataset, such as the user's favourite cable or over-the-air broadcasts could be used to predict when the TV is most likely in use. This dataset would need to be provided by the user using their Internet Movie Database (IMDB) profile or the TV Guide website.

Overall, a habits-based dataset would need to be active only when the user is on site. Tracking real-time geographic locations would be a good indicator of active device use versus maintenance or standby device usage. Devices that contain microcontrollers often turned on without user intervention to request software updates or perform routine maintenance or surveillance. Outdoor lighting, security cameras, sprinkler systems, and network-enabled devices would be examples in this category. By referencing a user's geographic location, an algorithm can determine if the user was not involved in the activation of the device and therefore the device was running autonomously. When multiple people use a particular location, knowing which particular user is on-site can help target the correct habit-based dataset.




Figure 8: On vs. off durations for common home appliances (figure from Hart, 1992)

Sequential events and usage trends can both contribute to a habits-based dataset. A sequential event takes place when the use of one device is followed by the use of a different device with a high probability. A disaggregation algorithm to improve choices when multiple matches exist for an event can use the correlation of use between multiple devices.

Usage trends result from use cases of devices. Some factors relating to a device use case are time of day, date (weekend, weekday, holiday), on duration, off duration, number of people onsite, and frequency of use. There are some devices that may turn on and off for a few minutes throughout the day.

Other devices may be left on in the morning, while others are typically used on weekends. Standard usage trends can be developed using typical use scenarios for locations by their function. For example, a washer and dryer would be used primarily on a weekend at any point during the day. In contrast to a washer, a dishwasher may be run daily in the evening. To further standard use trends, a residence would have a different profile to a call centre due to the tasks performed, type and number of devices on site, and number of people.

Using a sequence of events, the habits dataset could be enabled and disabled. For example, If using the coffee machine is followed by the garage door opening and closing, the habits dataset can be disabled. To improve the logic of toggling a habits dataset, devices that reveal a user active on-site can signal if the dataset should be enabled.

Models for presence detection could be made irrelevant if an alarm system's state is made available. A typical alarm system has three states; away, home, and disarmed. These three states can be simplified to two because the disarmed state does not provide any information about the users that are onsite. Ideally, we would use the energy signatures of an alarm system to indicate the system's state, but this is not possible because a system transition does not reveal a transient energy signature. Another input would be required to capture a system state change. Therefore, a network connected alarm system would be ideal for this occasion.

8.2. Conceptual Workflows for Context Awareness

Context can also be introduced through location awareness, of which there are 2 methods available. The first would require a user's location to be tracked, while the second would rely on device tagging, eliminating the need for a user's positional information. The goal of location awareness could be twofold. The first is to reveal a devices location so conservation efforts can be targeted to a specific device. The second is to allow the user to categorize or tag a device for future tracking and signature classification. In most cases, it is cost prohibitive to monitor individual devices because additional

hardware is required. The alternative is to minimize the impact by using the tracking capabilities of other systems. Through a user's positional information, devices the user interacts with can be noted. If device characteristics match changes to aggregate readings, it can be assumed the device is in use.

Geo-fencing or tracking a user's location in relation to a geographic boundary could be used to limit the number of signatures used for analysis of aggregate consumption. A catalogue of devices by location would need to be built to augment disaggregation with user location awareness. It is not necessary for a catalogue to be built in advance to implement disaggregation, instead a user could use a mobile application to tag and build the database of devices. A workflow involving the user in the system would enhance conservation efforts.

Integrating a location awareness database into a disaggregation workflow could allow servers that process data to run more efficiently and accurately. Knowing which subset of signatures to use in an initial search would reduce the pool of data. For signatures that are unknown, having user location information aids in the classification of an unknown signature. If the user's position were related to the use of a device, they would be able to recall and tag the signature. Having location awareness associated with aggregate data provides the opportunity to conserve energy by targeting specific devices.

Several technologies exist for implementing location awareness. The choice of which system to implement requires an assessment in regards to cost, accuracy relating location to device usage, user impact, complexity of integration, and ancillary benefits. This section will explore the methods used to track users within a building and focuses on their synergy with a disaggregation system. Although this paper does not contain an exhaustive list of location awareness technologies, it not only tries to cover systems that a business may already have present, but also focuses on emerging location awareness technologies and datasets.

<i>Table 3: Location awareness technologies</i>						
	Prevalence	Cost	Accuracy	Complexity	Impact	TTL
Alarm System	High	Low	Binary	Simple	Low	Short
Security Cameras	High	Medium	Varies	Moderate	Low	Short
Facial Recognition	Low	High	Varies	High	Low	Medium
Network Traffic	High	Low	Device	Low	Low	Short
WiFi Tracking	Low	Low	< 12cm	Moderate	Low	Medium
Swipe Cards	Medium	Medium	< 7m	Moderate	Medium	Short
Mobile Device	High	Low	< 300m	Low	High	Medium
Beacons	Low	High	~3cm	Moderate	Low	Short
Ultrasonic	Low	High	< 1cm	Moderate	Low	Short
Habit Tracking	Moderate	Low	Varies	High	High	Long
External Datasets	High	Low	Varies	High	Low	Short
Sub Metres	Low	Varies	Zone	Low	Low	Short
Socket Tagging	Low	High	Device	Low	High	Short

For some location integrations, the user's location may not be accurate to a particular device. Locations may relate to an area of the building, such as a zone or room. A disaggregation workflow would need to consider a user's proximity to a device rather than a particular room. To overcome this, a database would need to allow reference to a related set of devices. Creating a dataset of devices by location is tedious and may not be a viable long-term solution. For locations where devices are often moved, it may not be possible to maintain a catalogue of device locations.

8.2.1. Alarm Systems

A simple method to tracking a user's presence at a location is through an alarm system. These systems typically consist of a control panel, motion sensors, door contacts, window break sensors, siren, and a power supply. Newer systems provide smartphone applications for monitoring and control. Systems that provide API controls are ideal for integration with a disaggregation system because no modifications would be required. Low-end systems would require a control panel that is capable of sharing the system's status. If motion sensor information is shared while the system is not armed, it can be used to determine a user's location within a building. Since motion sensors are not available in all rooms, their accuracy is limited to spaces where they are installed. Therefore, they may not be an ideal candidate for location tracking within a building. Another downside to motion sensors is that they lack direction, so it is not clear if a user is working within, exiting, or entering an area.

Without access to motion sensors, security systems only provide onsite information for an entire location. If this information were made available, it would allow a system to dynamically select which devices are in use when people are onsite. Without unique control panel user codes, there is no distinction regarding who is onsite when the alarm is disabled, nor does it provide a user's location within the building. For a multi-user location, the lack of detail does not assist with disaggregation algorithms. It does work well for a residential scenario where residents are generally away during the day. There is also a potential issue when the system is not engaged, yet users are not onsite. Therefore, an alarm system has the potential to augment an algorithm, but it is beneficial to look at other integrations to provide deeper location awareness.

8.2.2. Security Cameras

Installing a camera system varies greatly in cost depending on the type of technology (IP or analogue). Regardless of the cost of individual cameras, the software to process camera feeds is where location awareness can be communicated. This takes place at a DVR (Digital Video Recorder), where

software is used to store and interpret footage. Majority of DVRs have motion tracking and remote access. The camera feeds can be captured using third party software to allow for further processing.

Tracking movement using a camera feed would indicate a user's location, but it does not reveal identity. Similar to motion sensors, security cameras are unable to easily associate position. Therefore algorithms would not implement individual user habits, rather bias detection towards the devices that are contained within a specific camera monitored locations within the building. In such a case, a disaggregation algorithm would need to associate a catalogue of devices to the location in which the motion was detected.

Facial recognition when paired with location would provide details of the individual. If devices can be located on camera feeds, their location can further dictate the devices they are interacting with. These details can be used to track habits of individual so their trends can also be used to target a subset of device signatures. Since the individual's details are being provided, an algorithm could then bias device usage if they have similar consumption characteristics.

The downside is the cost and infrastructure required. Processing video footage real-time can be resource intensive and expensive. Most commercial locations have security cameras within the premises, so it may be possible to integrate third-party software. The IT requirements to support such a system may be out of reach for most businesses. This paired with camera coverage within a location would make it an expensive undertaking.

8.2.3. Smartphone Location Tracking

Using a mobile device's location services would require a user to opt into the tracking system. In a commercial environment, it may be difficult to persuade employees to opt in to tracking. Unless mobile phones are owned and managed by the company, it may not be feasible. Therefore smartphone tracking would be ideal for residential users that are open to moving past privacy concerns to enable conservation

efforts. Aside from the installation of a mobile application, there is no additional infrastructure required to implement location based services (LBS).

LBS on a smartphone use a combination of cellular, WiFi and GPS to work. A-GPS (Assisted Global Positioning Systems) are generally employed by mobile devices because they contain don't contain a full GPS chip (Zendbergen, 2009). Instead a server-side proxy is used to interpret GPS signals and relayed to the mobile device. Beacons can also play a part in improving location accuracy and are covered separately in this section because they require additional infrastructure. WiFi is a ubiquitous network infrastructure and is remains a part of location services. There is an emerging technology that uses WiFi to sense movement in order to reveal location and implement gesture tracking. This technology is covered in *WiFi Tracking*.

The use of location based services on a mobile device varies over a kilometre at times. This is due to the poor cellular coverage and a lack of cellular towers for triangulation. To overcome this, a building can employ a dense mesh of wireless access points (AP). The positioning capabilities of multiple networks provides geolocation within 5m (Mok & Retscher, 2007). This would be ideal for a practical application for disaggregation but is expensive to employ over a large multi-floor building. For metropolitan locations, the Place Lab project was able to achieve positioning within 15 to 40m (Cheng et al, 2005).

<i>Table 4: Comparison of location services by median error</i>	
Type	Median Error
GPS	8m
WiFi	74m
Cellular (3G)	600m

8.2.4. Swipe Cards

In order to control entry points for security purposes, swipe card systems log a unique identification number for each user along with a zone ID that is associated with the swipe panel. These systems rely on a passive RFID tag or a magnetized strip that is embedded into a keycard. Implementing a swipe card system involved the cabling and mounting of panels throughout a premises. All panels communicate with a server that validates the keycard to a database containing all configured panels.

Typical RFID-based swipe panels employ a low-frequency communication allowing for up to 30cm ranges (Rieback et al, 2006). Systems that use higher frequencies in the 2.45GHz range allow for distances up to 7m. Longer distances could be used to actively check for an individual but cause door locks to open for other people if authorized users are in proximity to the sensor. An ideal case would be a panel that uses a low and high frequency. The low frequency communication would be used to authorize users, while the high frequency communication would be used for user tracking.

Since the user and location database are already part of a swipe card system's implementation, an appliance location database would need to be generated. Items in this additional database would relate location to the panel location IDs in the existing database. The subset of devices within the access area would then be given an increased probability of usage. Event detection algorithms would associate or tag events with access to a particular area.

The primary downside relates to entry tracking. Swipe cards are generally one-way systems, meaning only entry is being tracked. They do not enable the active location tracking of a user, nor do they explicitly enable knowledge on device usage. Instead, all information is inferred. If a user swipes their card, it is expected they are entering and actively participating within the space beyond the entry point. There are cases where a user may swipe and decide not to enter. For those users that enter a space, over time their location would grow stale and the probability they are remaining within the area decreases. Using a panel for all entry and exit would eliminate this at the cost of the user's ease-of-use.

8.2.5. WiFi Tracking

Routers with multiple antennae that support multiple-input and multiple-output (MIMO) transmissions have the ability to track the reflection of signals to track movement. Line of sight is not required for wireless frequencies to be used as a tracking medium. A single antenna when used has the ability to sense a change in object position, but does not discern direction. This is achieved using the properties of the Doppler effect, or a wave's change in frequency when the source moves relative to the observer (Gupta et al, 2012). When several antennae are used, object tracking in 3-space is enabled. As an added benefit, several moving objects can be tracked at once. WiFi tracking technologies would allow for disaggregation algorithms to benefit from user positioning as an input to improve accuracy.

Wireless tracking by Wi-Vi and WiSee rely on the 2.4GHz transmission frequency and MIMO antenna configurations. WiTrack uses a low power signal that sweeps from 5.46-7.25GHz every 2.5 milliseconds (Fadel et al, 2015). Although WiFi signals permeate solid walls, their accuracy is affected by building materials. Glass, wooden doors, and hollow walls cause a 1-way radio frequency attenuation between 3dB and 9dB (Cumberland, n.d.) . Concrete and steel have a greater impact, fairing over 18dB and 40dB respectively. Considering materials in use at a location employing this technology would be an important consideration.

Gesture tracking is also enabled using Wi-Vi frequency tracking. Using the systems gesture tracking accuracy of 100% within 8m of the antennae, human-device interactions can be tracked (Adib et al, 2013). The WiSee system has an average accuracy of 95% using a single transmitter and four receivers (Pu et al, 2013). Setting this up would require training for every device. If a device is ever relocated, training would need to be performed again. Training could be an automated process if change events were linked with user locations and gestures. Over time, validating events by location could be used for autonomous tagging. Tags would allow a user to classify the event asynchronously. This would then be used to link an action to a particular device signature.

The three WiFi tracking systems mentioned in this section have not revealed the accuracy implications of a wireless mesh network spanning large areas. Without further testing, it is unknown if using 802.11 wireless signals are viable for a large facility. For smaller locations, using wireless signals to track movement and position may hold the key to location tracking. More research needs to be conducted on tracking users in multiple rooms using broadly located antennae or APs.

8.2.6. Ultrasonic

Ultrasonic systems are highly accurate for location tracking. Often placed in the ceiling, a hardware device known as a Mobile Positioning Device, or MPD, emits a pulse. The MPD or another receiver device calculates the pulse times-of-flight. Receivers are placed at known locations and relay the pulse measurements to a computing device known as a matrix manager. Through a process known as multilateration, user and object positions are revealed. The entire process is similar to a bat or dolphins' navigation system.

Ultrasonic systems such as the Active Bat can pulse up to 50 times per second. The positioning information it provides is accurate within 3cm 95% of the time (Madhavapeddy & Tse, 2005). Such systems would allow location data to be matched to device usage within a location, providing a matrix of MPDs and receivers is installed throughout the location.

Similar to beacons, an ultrasonic system requires the installation and calibration of hardware. Although the unit price of receivers is low, the cost to run cables to a matrix of receivers over a large area would be justifiable simply for conservation efforts.

8.2.7. Beacons

Bluetooth enabled beacons require a user to carry a mobile device to detect position. The cost to integrate beacons is reliant on the square-footage of a building. Each beacon must be placed at location intervals that are determined by the manufacturers specifications. The detection of an individual

smartphone requires a process known as fingerprinting. Through this process, a user's location can be determined to 3cm in 3D at 95% confidence (Fangher et al, 2014).

The energy use of a beacon is relative low compared to WiFi. Therefore the impact to the end user is low while offering improved accuracy over a mobile phone's location services. Installing beacons usually requires cabling to be run to each deployment. Wireless beacons are also an option, but replacing batteries is tedious.

8.2.8. Network Traffic

At a corporate and residential level, a gateway can be used to analyze network traffic. Cross-referencing network data with device consumption information can be used to validate the device in use through its media access control address (MAC address). Since a MAC address is a unique identifier, it provides information on the manufacturer and can be traced back to an individual device. This match is ideal for network enabled devices such as computers, smart TVs, printers, and network hardware. As more devices are network enabled, it will be possible to track the usage of lights, kitchen appliances, electrical sockets, and more.

Accuracy in the case of network traffic may be impacted by the delay in communication. Network interfaces often take time to come online when a system is turned on. A disaggregation system would need to take this into account and track the duration a MAC address is online relative to the aggregate consumption edge detection. In the case of a desktop computer booting up, it may take over a minute for a network interface controller (NIC) to come online, while shutdown may take place immediately. It would be necessary for a disaggregation algorithm to account for temporal variations in network traffic relative to electrical consumption.

For devices left on standby, their NIC may be online the entire time to allow for wake-on-LAN protocols or scheduled updates. In these cases, the number of frames being transferred to the device would reveal its state. Since Ethernet frames are sent with MAC address information within the headers, it

would not be difficult to track active usage versus standby usage. However, tracking standby usage can also be used to reveal vampire power. The actual magnitude of a device's consumption in standby mode may not be revealed without sub-second readings. However, it may be noted that the device is an active contributor to the noise floor.

To integrate a disaggregation system with a network traffic analysis would require a router or firewall to be configured to share a portion of layer 2 communications. The raw network data would not be needed, just a portion of frame headers. An Ethernet frame has 8 parts consisting of a preamble, start frame delimiter, destination and source MAC address, frame length, data, padding, and checksum (Mackenzie, 2011). The only item required is the source MAC address originating from the device along with a timestamp. An open firmware could be modified for this task or a piece of hardware installed at a cost of under \$100. The impact to the end user would be negligible.

9. Conclusion

9.1. Summary of Findings

Accurately disaggregating energy through NILM has the potential to improve our conservation efforts. The adoption of smart metres in homes and businesses produces a timely opportunity to harness the data being produced. The downside is that sampling rates of smart metres are not detailed enough for the use of disaggregation of more than 8 appliances. Unless hardware and firmware changes take place, the data produced will only allow for lighting and mid-size devices. In cases where high-efficiency appliances have already been installed, the next step would be to tackle moderate and low powered devices that significantly contribute to consumption in both residential and business environments. Small electric devices, heating elements, and motors account for 27% of residential consumption. When paired with lighting, entertainment systems, and computers, the total is 47% (UEIA, 2015).

Advances in disaggregation algorithms continue due to work being done in image and audio source separation. There is no doubt that accuracy will improve but unless context can be provided, it may be unfeasible to expect consistency. Either sampling rates need to exceed 60Hz, or external datasets need to be included. APIs can be used to access datasets, but efforts need to be made regarding the integrating and development of logic to support this addition.

From the research conducted, it appears CFHSMM is an algorithm suited for disaggregation. Its allowance for external contextual datasets while keeping computing requirements low is ideal for disaggregation. Coupling context with a crowd sourced database of device signatures would enable a large number of devices to be detected without engaging the user in a training process. Another consideration is the low sample rate available with smart metres. If sample rates are increased to allow for harmonic analysis, it may not be necessary to include contextual datasets. The detail and accuracy offered by algorithms processing high-frequency samples would require a shared signatures database and device locations. Having this information would improve classification and targeted conservation efforts.

9.2. Future Considerations/Research

A practical application of crowdsourcing signatures for disaggregation should be conducted. Determining quorum, number of false positives, database growth rate, and creating a database of signatures would further research disaggregation development. It would require more people to adopt open monitoring systems that allow for changes to hardware and software configurations.

Since smart metres are being provided to customers at low costs, the customer is denied access to modifying the hardware. Until utility providers see the value in higher sample rates, it will be up to the customer to find a work around. The only viable alternative to smart metres for NILM is the induction clamp. Pairing clamps with software through the open-source community and private sector enables the general user to access the power of disaggregation. The process is not as easy as it should be, but there have been many advances in this aspect.

There is still the opportunity to refine the accuracy of classification algorithms such as ANN, DTW, SVM, k-NN, and FHMM. The situations in which one algorithm should be used over another would reveal no single combination of classification, training, and event detection workflows will work as a one-size-fits-all solution. Instead, testing needs to be done to dynamically select algorithms according to a location's profile. It may be that the device generating data will dictate the ideal workflow to use.

The increasing adoption of network-enabled devices may allow for tagging or external datasets to be generated. It will take time for this class of devices to work with a disaggregation system. The question of interoperability would be paramount to having a fully integrated disaggregation system. Providing an standard is created and adopted by manufacturers, network enabled devices may come with the added benefit of remote control. Chaining devices together would allow for area or process driven automation.

Technological breakthroughs and the penetration of energy monitoring systems may lead to a large improvement in conservation efforts. The improved efficiency in appliances coupled with users changing their habits will help lower usage, but history reveals there is often an opposite effect. Even with lowering resource usage, the growing population and demand continues to increase. Jevon's paradox described the soaring consumption of coal in England in the 19th century (Alcott, 2005). Due to the efficiencies created by James Watt's coal-fired steam engine, travel became more affordable. People travel more often and further than ever before. Through the replacement of equipment with energy efficient equivalent, the paradox is that their usage increases and we as a global population install more products.

10. Bibliography of Sources & Materials

Adib, F., & Katabi, D. (2013). See through walls with WiFi! (Vol. 43, No. 4, pp. 75-86). ACM.

Adib, F., Kabelac, Z., & Katabi, D. (2015, May). Multi-person localization via RF body reflections. In *Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation* (pp. 279-292). USENIX Association.

Alcott, B. (2005). Jevons' paradox. *Ecological economics*, 54(1), 9-21.

Anderson, C. (2008). *The long tail: Why the future of business is selling less of more* (Rev. and updated ed.). New York: Hyperion.

Armel, K., Gupta, A., Shrimali, G., & Albert, A. (2013). Is disaggregation the holy grail of energy efficiency? The case of electricity. *Energy Policy*, 52, 213-234.

Belkin (n.d.). Belkin : Conserve : Socket : Saving money and energy is easy with Conserve Socket. Retrieved August 3, 2016 from <https://www.belkin.com/conserve/socket/>

Bronkhorst, A. W. (2000). The cocktail party phenomenon: A review of research on speech intelligibility in multiple-talker conditions. *Acta Acustica united with Acustica*, 86(1), 117-128.

Chang, H. H., Lin, C. L., & Lee, J. K. (2010, April). Load identification in nonintrusive load monitoring using steady-state and turn-on transient energy algorithms. In *Computer supported cooperative work in design (cscwd), 2010 14th international conference on* (pp. 27-32). IEEE.

Cheng Y., Chawathe Y., LaMarca A., and Krumm J. (2005) Accuracy Characterization for Metropolitan-scale Wi-Fi Localization. Seattle, WA, Intel Research Report No. IRS-TR-05-003

City of Cumberland (n.d.), How Signal is affected, City of Cumberland Report. Retrieved from www.ci.cumberland.md.us

- Darby, S. (2006). The effectiveness of feedback on energy consumption. Technical report, Environmental Change Institute, University of Oxford.
- Elafoudi, G., Stankovic, L., & Stankovic, V. (2014, May). Power disaggregation of domestic smart meter readings using dynamic time warping. In *Communications, Control and Signal Processing (ISCCSP), 2014 6th International Symposium on* (pp. 36-39). IEEE.
- Elhamifar, E., & Sastry, S. (2015, October). Energy Disaggregation via Learning ‘Powerlets’ and Sparse Coding. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.
- Faragher, R., Harle, R. (2014) An Analysis of the Accuracy of Bluetooth Low Energy for Indoor Positioning Applications.
- Faruqui, A., Sergici, S., & Sharif, A. (2010). The impact of informational feedback on energy consumption—A survey of the experimental evidence. *Energy*, 35(4), 1598-1608.
- Figueiredo, M., De Almeida, A., & Ribeiro, B. (2012). Home electrical signal disaggregation for non-intrusive load monitoring (NILM) systems. *Neurocomputing*, 96, 66-73.
- Ghahramani, Z., & Jordan, M. I. (1997). Factorial hidden Markov models. *Machine learning*, 29(2-3), 245-273.
- Gonçalves, H., Ocleanu, A., Pandya, A., & Tragut, M. (2010) Non-Intrusive Load Monitoring.
- Gupta, S., Morris, D., Patel, S., & Tan, D. (2012, May). Soundwave: using the doppler effect to sense gestures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1911-1914). ACM.
- Hart, G. W. (1984). Non-intrusive Appliance Load Data Acquisition Method. MIT Energy Laboratory Technical Report.

Kelly, D., & Knottenbelt, W. (2011). Disaggregating Smart Meter Readings using Device Signatures.

Imperial Computing Science MSc Individual Project.

Kelly, J. (2014, April 13). Nonintrusive Load Monitoring Toolkit (NILMTK). Retrieved March 6, 2015,

from <https://github.com/nilmk/nilmk>

Kim, H., Marwah, M., Arlitt, M. F., Lyon, G., & Han, J. (2011, April). Unsupervised Disaggregation of

Low Frequency Power Measurements. In *SDM* (Vol. 11, pp. 747-758).

Kolokotsa, D., Rovas, D., Kosmatopoulos, E., & Kalaitzakis, K. (2011). A roadmap towards intelligent

net zero-and positive-energy buildings. *Solar Energy*, 85(12), 3067-3084.

Kolter, J. Z., & Johnson, M. J. (2011, August). REDD: A public data set for energy disaggregation

research. In *Workshop on Data Mining Applications in Sustainability (SIGKDD)*, San Diego, CA

(Vol. 25, pp. 59-62).

Kolter, J. Z., & Jaakkola, T. (2012). Approximate inference in additive factorial hmms with application to

energy disaggregation. In *International conference on artificial intelligence and statistics* (pp.

1472-1482).

Kolter, J. Z., Batra, S., & Ng, A. Y. (2010). Energy disaggregation via discriminative sparse coding. In

Advances in Neural Information Processing Systems (pp. 1153-1161).

Kouemou, G. L. (2011). *History and theoretical basics of hidden Markov Models*. INTECH Open Access

Publisher.

Kyan, M. (2015, August 26). [Online interview].

Lee, W. L., & Burnett, J. (2008). Benchmarking energy use assessment of HK-BEAM, BREEAM and

LEED. *Building and Environment*, 43(11), 1882-1891.

- Ling, P. J., & Eldridge, C. J. (1994, September). Designing modern electrical systems with transformers that inherently reduce harmonic distortion in a PC-rich environment. In *Proc. Power Quality Conf* (pp. 166-178).
- Lisovich, M., & Wicker, S. (2008, March). Privacy concerns in upcoming residential and commercial demand-response systems. In *Proc. of the Clemson University Power Systems Conference (Clemson, SC*.
- Marsland, S. (2014). Machine learning: an algorithmic perspective. CRC press.
- McLaughlin, S., McDaniel, P., & Aiello, W. (2011, October). Protecting consumer privacy from electric load monitoring. In Proceedings of the 18th ACM conference on Computer and communications security (pp. 87-98). ACM.
- Mackenzie, L. (2011). Ethernet Frame Format. Retrieved August 3, 2015.
- Makonin, S. (2012). Approaches to Non-Intrusive Load Monitoring (NILM) in the Home.
- Madhavapeddy, A., & Tse, A. (2005). A study of bluetooth propagation using accurate indoor location mapping. In UbiComp 2005: Ubiquitous Computing (pp. 105-122). Springer Berlin Heidelberg.
- Mercier, C, & Moorefield, L., 2011. Commercial office plug load savings and assessment: Final report. Produced by ECOVA and supported through the California Energy Commission's Public Interest Energy Research Program. Available at http://www.efficientproducts.org/reports/plugload/Ecova_Office%20Plug%20Load%20Assessment_Final%20Report%20Nov%202011.pdf.
- Mok, E. and Retscher, G. (2007) Location determination using WiFi fingerprinting versus WiFi trilateration. Journal of Location Based Services 1: 145–59
- Navigant Research (2014, August 27). Smart Meters. Retrieved from <https://www.navigantresearch.com/research/smart-meters>

P3 International. (n.d.). Kill A Watt Meter - Electricity Usage Monitor. Retrieved June 29, 2015 from <http://www.p3international.com/products/p4400.html>.

Peaksaver PLUS. (n.d.). Retrieved August 28, 2015 from <https://saveonenergy.ca/Consumer/Programs/PeaksaverPlus.aspx>.

Piotrowski, N. (n.d.). True Power vs. Apparent Power: Understanding the Difference. Retrieved August 10, 2015 from <http://www.aptsources.com/resources/pdf/True%20vs.%20Apparent%20Power.pdf>

Pu, Q., Gupta, S., Gollakota, S., & Patel, S. (2013, September). Whole-home gesture recognition using wireless signals. In *Proceedings of the 19th annual international conference on Mobile computing & networking* (pp. 27-38). ACM.

Peterson, (2007). International Journal of Sustainability in Higher Education, Vol. 8 No. 1, 2007, pp. 16-33

Rieback, M. R., Crispo, B., & Tanenbaum, A. S. (2006). The evolution of RFID security. *IEEE Pervasive Computing*, (1), 62-69.

Sakoe, H., & Chiba, S. (1978). Dynamic programming algorithm optimization for spoken word recognition. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 26(1), 43-49.

Sankaran, C. (1999). Electrical Construction and Maintenance Magazine. Retrieved August 3, 2015 from <http://ecmweb.com/power-quality/effects-harmonics-power-systems>

Siau, N. (2004). Power harmonics analysis for electrical device signature identification using the artificial neural network and support vector machine. University of Singapore, Open (Master's Theses).

Silva, C. (2005). Variable-Inductance Transducers. In *Mechatronics: An integrated approach*. Boca Raton: CRC Press.

SmartHome (n.d.). Smarthome Home Automation Superstore. Retried August 3, 2016 from

<http://www.smarthome.com/>.

Snow, R., O'Connor, B., Jurafsky, D., & Ng, A. Y. (2008, October). Cheap and fast---but is it good?: evaluating non-expert annotations for natural language tasks. In Proceedings of the conference on empirical methods in natural language processing (pp. 254-263). Association for Computational Linguistics.

TED (2012). TED 5000 Third Party Posting API (The Energy Detective Corporation). Retrieved from

<https://www.theenergydetective.com/downloads/tedproapi.pdf>

UEIA (2015). Annual Energy Outlook 2015 (Table 4). U.S. Energy Information Administration.

Ward, A. M. R. (1998). Sensor driven computing (Doctoral dissertation, University of Cambridge).

Xu, Y., & Milanovic, J. V. (2015). Artificial-intelligence-based methodology for load disaggregation at bulk supply point. *Power Systems, IEEE Transactions on*, 30(2), 795-803.

Xudong, Y., Jingping, J., Brierley, S. D., Chiasson, J. N., Lee, E. B., Zak, S. H., ... & Polycarpou, M. M. (1998). Adaptive nonlinear design without a priori knowledge of control directions. Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.33.5444>.

Zandbergen, P. A. (2009). Accuracy of iPhone locations: A comparison of assisted GPS, WiFi and cellular positioning. *Transactions in GIS*, 13(s1), 5-25.