

Transferring Decomposed Tensors for Scalable Energy Breakdown across Regions

Nipun Batra, Yiling Jia, Hongning Wang, Kamin Whitehouse
University of Virginia

Abstract

Homes constitute roughly one-third of the total energy usage worldwide. Providing an *energy breakdown – energy consumption per appliance*, can help save up to 15% energy. Given the vast differences in energy consumption patterns across different regions, existing energy breakdown solutions require instrumentation and model training for each geographical region, which is prohibitively expensive and limits the scalability. In this paper, we propose a novel region independent energy breakdown model via statistical transfer learning. Our key intuition is that the *heterogeneity* in homes and weather across different regions most significantly impacts the energy consumption across regions; and if we can factor out such heterogeneity, we can learn region independent models or the *homogeneous* energy breakdown components for each individual appliance. Thus, the model learnt in one region can be *transferred* to another region. We evaluate our approach on two U.S. cities having distinct weather from a publicly available dataset. We find that our approach gives better energy breakdown estimates requiring the least amount of instrumented homes from the target region, when compared to the state-of-the-art.

Introduction

Homes account for roughly one-third of the total energy consumption worldwide (Pérez-Lombard, Ortiz, and Pout 2008). Previous research has demonstrated that providing an *energy breakdown*: itemising total energy into individual appliances, such as lighting, heating, etc., can help occupants save up to 15% energy (Darby 2006; Armel et al. 2013). Previous research has also discussed various other benefits of energy breakdown, such as load forecasting, understanding appliance usage for improved design, policy making, among others (Armel et al. 2013; Batra, Singh, and Whitehouse 2015).

Various methods for providing an energy breakdown have been studied. The most intuitive way to get an energy breakdown involves instrumenting each appliance with a sensor. Various sensing systems have been proposed in the past (DeBruin et al. 2015; Jiang et al. 2009). However, these sensing systems require extensive installation and are thus prohibitively expensive to scale across a large number of homes.

In contrast, since the 1980s a novel technique called non-intrusive load monitoring (NILM) has been proposed, which uses statistical techniques to break down the energy measured at the home meter level (Hart 1992). However, even NILM would require installing a sensor (such as a smart meter) and would thus cost up to \$500 per home, limiting the scalability. Recently, there have been works (Batra et al. 2017; Batra, Singh, and Whitehouse 2016) on providing an energy breakdown without any hardware installation. These approaches can provide an energy breakdown just using the monthly electricity bills and a few homes in the region which already have an energy breakdown. The approaches promise considerable improvement in scalability, but they impose strong assumptions about the problem: the training homes are similar to the testing homes (i.e., all homes are identical and independently distributed). As the set of training homes grows and begins to span multiple climate zones and varied homes (old v/s new, well v/s poorly insulated, studio v/s 3 BHK, etc.), the error will inevitably increase.

In this paper, we present a novel region independent energy breakdown method. Our key insight is that the *heterogeneity* in homes and weather across different regions most significantly impacts the energy consumption across regions; and if we can factor out the weather and homes, we can learn region independent models or the *homogeneous* energy breakdown components for the appliances. As previously shown (Batra et al. 2017; Batra, Singh, and Whitehouse 2017), the heterogeneity across homes (e.g., well-insulated v/s poorly insulated homes) can be captured using a low dimension representation. Intuitively, homes form clusters in this low dimension space. Similarly, the energy dependence of different appliances concerning weather (e.g., the cooling load may be directly proportional to temperature, fridge energy may be season independent, etc.) can also be encoded using a low-dimensional representation. Once we account for the heterogeneity, we can learn *homogeneous* or region-independent factors about the appliances, which capture the interaction between homes and weather. An example of such a factor would be the direct relationship of cooling on home insulation and external temperature. Another example could be the weak relation between fridge energy and home insulation and external temperature.

Our approach thus boils down to learning the heterogeneous home and season/weather factors, referred to as **H**

and $\mathbf{S}$ respectively; and learning the appliance factor ($\mathbf{A}$) which depends on $\mathbf{H}$ and $\mathbf{S}$. More specifically, we assume the appliance factor $\mathbf{A}$ itself is a three-way tensor (spanning over the number of appliances, and the dimension of $\mathbf{H}$ and $\mathbf{S}$). Intuitively, each appliance factor can be considered as a set of linear combinations of season factors, which form a set of *home bases* (e.g., appliances whose energy usage is sensitive to the change of weather v/s those insensitive). Moreover, each home can thus be characterised as a linear combination over those home bases (e.g., well v/s poorly insulated homes). Therefore, this cross region energy breakdown problem has been naturally formalised as a tensor factorisation problem, where appliance factor $\mathbf{A}$ could be reused across regions. Distinct from standard tensor decomposition solutions (such as PARAFAC (Harshman 1970) and Tucker (Tucker 1966)), which assume the decomposed factors are independent, our solution explicitly encodes the inter-dependency pattern between home factor $\mathbf{H}$ and season factor $\mathbf{S}$ within regions via the appliance factor $\mathbf{A}$, and therefore has the potential to better factor out region dependence. We evaluate our approach on a publicly available dataset called Dataport (Parson et al. 2015). We learn $\mathbf{A}$ factors from 534 homes in Austin and used 40 testing homes from San Diego to evaluate the prediction energy breakdown. Our approach gives better accuracy compared to five baseline approaches, in particular for a low amount of adaptation data required. We see a similar trend when we perform a transfer of $\mathbf{A}$ from 39 homes in San Diego to 40 homes in Austin.

Related Work

Since George Hart’s seminal work on non-intrusive load monitoring (NILM) in the early 1980s (Hart 1992), the research community has proposed several solutions to scale up energy breakdown. Primarily, the work can be categorised into: 1) metering hardware, and 2) NILM algorithms. It should be noted that both these lines of work require instrumentation across each home.

Various metering hardware based approaches for energy breakdown have been proposed in the past. A few of these involve directly instrumenting an appliance with a power sensor (DeBruin et al. 2015). Many indirect load sensing approaches have also been raised in the past, which monitor the power of an appliance via some proxy signal. For example, using magnetic field as a proxy for power (Kim et al. 2009), or EMI noise generated by power supplies as a proxy for the power consumption of electronic appliances (Gupta, Reynolds, and Patel 2010).

NILM algorithms work on the premise that we have to perform source separation on the power signal measured at a single point (home mains). Various NILM algorithms have been proposed in the past three decades (Zoha et al. 2012; Armel et al. 2013). Majority of these algorithms work on time-series data obtained from a smart meter, collected at rates from 10s of kHz to a reading once every 15 minutes (Parson et al. 2012; Kolter, Batra, and Ng 2010; Kolter and Jaakkola 2012; Shao, Marwah, and Ramakrishnan 2013). To the best of our knowledge, these techniques have thus far only been illustrated on homes from the same

region, and of course, require metering hardware in each home, and thus are not scalable across regions.

In contrast, there has been some recent work (Batra et al. 2017; Batra, Singh, and Whitehouse 2016) which does not require any hardware to be installed in a test home of interest. The key idea behind such work is that “similar homes would have a similar per-appliance energy consumption”. These approaches could estimate the energy breakdown of a home by finding a similar home (based on monthly bills) that already has an energy breakdown available. Not only were these approaches shown to be more scalable, by requiring instrumentation in only a small set of homes, they were also shown to be more accurate compared to the state-of-the-art NILM approaches. However, such approaches have a fundamental issue – they promise huge improvement in scalability but require that the training homes be similar to the testing homes. As the set of training homes grow and begin to span multiple climate zones, the error will increase. In contrast, our approach aims to learn region-independent models which factor out the weather and the home differences across regions.

We now discuss the related work in the transfer learning domain. Transfer learning refers to a framework of statistical machine learning methods, which aim at reusing knowledge gained while solving one problem and to different but related problems. Various types of transfer learning solutions have been proposed, including model-based transfer (Bonilla, Chai, and Williams 2008), feature-based transfer (Raina et al. 2007), and instance-based transfer (Dai et al. 2007). A comprehensive survey can be found at (Pan and Yang 2010). Our solution can be considered as a model-based transfer learning approach, in which we estimate and reuse the learnt appliance factor $\mathbf{A}$ across regions. Both heterogeneities across homes and weather and homogeneity across regions are captured, because of our carefully imposed tensor structure. To the best of our knowledge, this is the first work applying transfer learning ideas in tensor decomposition, not just related to the specific energy breakdown problem.

Problem Statement

Our aim is to use energy data from a *source* region ($\mathbf{E}^{\text{source}}$) and energy data from a small number of *adaptation* homes from the *target* region, to estimate the energy breakdown across homes of the target region. We formally define our *energy tensor* ($\mathbf{E}_{\mathbf{M} \times \mathbf{N} \times \mathbf{T}}$) as a 3-way tensor where the cells contain energy readings of $\mathbf{M}$ homes for $\mathbf{N}$ appliances for $\mathbf{T}$ months. We consider household aggregate energy as one of the “special” appliances. Since aggregate data is easy to collect (via monthly electricity bills), we consider it to be always observed.

Thus, our problem statement can be generally formalised as: given source region energy tensor $\mathbf{E}_{\mathbf{M} \times \mathbf{N} \times \mathbf{T}}^{\text{source}}$ and target region energy tensor from a small set of adaptation homes $\mathbf{E}_{\text{Adapt} \times \mathbf{N} \times \mathbf{T}}^{\text{target}}$, we want to complete the energy tensor for test homes in target region $\mathbf{E}_{\text{Test} \times \mathbf{N} \times \mathbf{T}}^{\text{target}}$. It should be noted that even in the test region, we have aggregate energy always available ($\mathbf{E}_{\text{Test} \times 1 \times \mathbf{T}}^{\text{target}}$)

Approach: Transferable Tensor Factorisation (TTF)

Our core intuition is that the *heterogeneity* in homes and weather across different regions most significantly impacts the energy consumption across regions; and if we can factor out the weather and homes, we can learn region independent models or the *homogeneous* energy breakdown components. Our energy tensor has three dimensions - home, appliance and weather/season. Since homes and weather are inherently heterogeneous, we would want the appliance dimension to capture the homogeneity across regions.

Previous work (Batra et al. 2017) has shown that we can represent the heterogeneity across homes using a low dimension representation. Examples of such low-dimensional representation could be the home insulation, or the number of occupants, the area of the home, etc. Similarly, the energy dependence of different appliances with respect to weather/season (e.g. cooling load may be directly proportional to temperature, fridge energy may be season independent, etc.) can also be encoded using a low-dimensional representation. Thus, on the lines of previous work (Batra et al. 2017), we can decompose the energy tensor into three factors: home factors ($\mathbf{H}$), appliance factors ($\mathbf{A}$), and season factors ($\mathbf{S}$). We call this tensor decomposition structure as standard tensor factorisation (STF), where, each of $\mathbf{H}$, $\mathbf{A}$ and $\mathbf{S}$ are independent matrices. In STF, the dimensions of $\mathbf{H}$, $\mathbf{A}$, $\mathbf{S}$ are: $M \times r$, $N \times r$, and $T \times r$, where r is the rank of the energy breakdown tensor.

A fundamental issue with STF is that it does not explicitly factor out the heterogeneity across homes and seasons, as it assumes all three factors are independent from each other. Thus, there is no guarantee that which factor should account for the homogeneity across regions; though it may do a reasonable job in energy breakdown for a single region. To address this problem, we introduce our approach – Transferable Tensor Factorisation (TTF). As mentioned in our assumption, $\mathbf{A}$ should be learnt as a region-independent factor. Thus, we modify $\mathbf{A}$ to be another three-way tensor (spanning over the number of appliances, and the dimension of $\mathbf{H}$ and $\mathbf{S}$). Intuitively, each appliance factor can be considered as a set of linear combinations of season factors, which form a set of *home bases* (e.g., appliances whose energy usage is sensitive to the change of weather v/s those insensitive). And each home can thus be characterised as a linear combination of those home bases (e.g., well v/s poorly insulated homes). Therefore, this cross region energy breakdown problem has been naturally formalised as a tensor factorisation problem, where appliance factor $\mathbf{A}$ could be reused across regions. A caveat of TTF is that it requires more parameters to be learnt compared to STF; which is the trade-off we need to make for learning a region-independent $\mathbf{A}$ factor.

The key idea of having different heterogeneous and homogeneous factors in energy breakdown is that we only need to learn the heterogeneous components in a new region since the homogeneous components are region independent. This greatly reduces the amount of training data needed from a new region. Based on this idea, our overall procedure for estimating energy breakdown in a target region consists of

two steps. In the first step called **Normal learning**, we learn $\mathbf{H}$, $\mathbf{A}$ and $\mathbf{S}$ from the source region using $\mathbf{E}_{M \times N \times T}^{\text{source}}$. In the second step called **Transfer learning**, we reuse the $\mathbf{A}$ factor learnt from a source domain and only need to learn the $\mathbf{H}$ and the $\mathbf{S}$ factors from a few adapt homes. We now describe the two steps.

Normal Learning

In normal learning, we learn a model for energy breakdown in a given region. Our idea is to decompose the Energy Breakdown Tensor ($\mathbf{E}_{M \times N \times T}$) into three factors: i) Home factor ($\mathbf{H}_{M \times h}$), ii) Appliance factor ($\mathbf{A}_{h \times N \times s}$), and iii) Season factor ($\mathbf{S}_{s \times T}$), where h and s represent the number of home and season factors, respectively.

The normal learning energy tensor decomposition can be represented as:

$$\text{Min } \|\mathbf{E} - \mathbf{HAS}\|_{\mathbf{F}}^2 \text{ s.t. } \mathbf{H}, \mathbf{A}, \mathbf{S} \geq 0 \quad (1)$$

It should be noted that we enforce non-negativity constraints on each of $\mathbf{H}$, $\mathbf{A}$, and $\mathbf{S}$. This is because energy is a non-negative quantity; hence, each of $\mathbf{H}$, $\mathbf{A}$, and $\mathbf{S}$ can only non-negatively contribute to the overall energy.

Transfer Learning

For our transfer learning approach, we assume that the appliance factor ($\mathbf{A}$) is directly transferable across regions. Thus, for a given target region, we only need to learn the home ($\mathbf{H}_{\text{target}}$) and the season factors ($\mathbf{S}_{\text{target}}$), which are learnt as per the following optimisation problem,

$$\text{Min } \|\mathbf{E}_{\text{target}} - \mathbf{H}_{\text{target}} \mathbf{A}_{\text{source}} \mathbf{S}_{\text{target}}\|_{\mathbf{F}}^2 \text{ s.t. } \mathbf{H}_{\text{target}}, \mathbf{S}_{\text{target}} \geq 0 \quad (2)$$

If the appliance breakdown of the source domain is representative of the target domain, TTF is expected to work better than normal learning. Otherwise, as more training homes become available, normal learning will eventually do better than transfer learning as it learns an appliance breakdown that is more specific to the target domain.

We can solve the optimisation problem of $\text{Min} \|\mathbf{E} - \mathbf{HAS}\|$ using gradient descent. However, as these two optimisation problems are of high dimension and nonconvex, vanilla gradient descent would easily suffer from local minimums. However, variations of gradient descent such as Adagrad (Duchi, Hazan, and Singer 2011) have been shown to converge quicker and also use a different learning rate per dimension. Thus, we choose to use Adagrad for updating $\mathbf{H}$, $\mathbf{A}$, and $\mathbf{S}$. We use Autograd (Maclaurin, Duvenaud, and Adams 2015) for numerical gradient computation. We use the concept of projected gradient descent (Lin 2007) to ensure that $\mathbf{H}$, $\mathbf{A}$, and $\mathbf{S}$ are nonnegative. The following equation shows the procedure of projected gradient descent applied to a variable X . It would apply similarly to update $\mathbf{H}$, $\mathbf{A}$ and $\mathbf{S}$.

$$X_{i+1} = \begin{cases} X_i - \delta X_i \times \eta_i, & \text{if } X_i - \delta X_i \times \eta_i \geq 0 \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where X_i , δX_i , η_i represent the value of X , gradient of X , and learning rate at i^{th} iteration.

	HVAC	Fridge	MW	DW	WM	Oven
Austin	0.20	0.09	0.01	0.01	0.01	0.02
San Diego	0.12	0.15	0.02	0.02	0.01	0.03

Table 1: Proportion of energy consumed by different appliances across Austin and San Diego for the year 2014.

* In the summer months, the HVAC and the fridge can together account for $\approx 70\%$ of aggregate across both regions.

** We use the following abbreviations for appliances: Dish-washer (DW), Microwave (MW), Washing machine (WM)

Evaluation

Dataset

In this paper, we use the Dataport (Parson et al. 2015) dataset for evaluation. While Dataport contains data from various cities in the USA, the maximum instrumentation exists in Austin (Texas) with 534 homes, San Diego (California) with 39 homes, and Boulder (Colorado) with 40 homes. For Boulder, a significant amount of energy consumption comes from appliances which do not have observations in San Diego or Austin. More details about this limitation can be found in the Limitations section. Thus, in this paper, we only consider Austin and San Diego. Energy data from each appliance and the household mains (aggregate) was sampled every minute. Since we plan to evaluate the approach when only monthly electricity data in the form of bills are available, we downsample the data to a month resolution.

Figure 1 shows the energy breakdown across the two regions. It can be seen that Austin has a higher energy footprint compared to San Diego. This can be highly attributed to weather. Austin is a warmer city and thus has a higher cooling energy requirement. Table 1 shows the proportion of energy contributed by different appliances across the two regions. It is interesting to note that in Austin, HVAC contributes more compared to the fridge; however, in San Diego, the vice-versa trend is visible. The differences in absolute scales in energy consumption, differences in weather, differences in appliance contribution across the two cities, make our problem of transfer learning challenging and interesting.

Baselines

We compare our approach under normal and transfer settings with two approaches. First, we compare against the state-of-the-art approach that leveraged a matrix factorisation (MF) decomposition (Batra et al. 2017). The MF approach was only formulated for a single region, which we call as MF under normal settings. In the $\langle \text{MF}, \text{Normal} \rangle$ approach, for each appliance w , a matrix $X_w \in R^{M \times (2 \times T)}$ is created. The first T columns in this matrix represent the aggregate energy consumption and the last T columns represent the energy consumption of the w^{th} appliance. In this approach, the decomposition for each X_w is done as follows:

$$\begin{aligned} \text{Min } \|\mathbf{X}_w - \mathbf{Y}_w \mathbf{Z}_w\|_F^2 \\ \text{s.t. } \mathbf{Y}_w, \mathbf{Z}_w \geq 0 \end{aligned} \quad (4)$$

where $\mathbf{Y}_w$ and $\mathbf{Z}_w$ correspond to the latent factors for homes and $\langle \text{appliance}, \text{months} \rangle$ respectively. For transfer learning ($\langle \text{MF}, \text{Transfer} \rangle$), we use the $\mathbf{Z}_w$ learnt from the source domain and learn $\mathbf{Y}_w$ from the target domain. Since MF was shown to be better than the state-of-the-art NILM approaches, we do not compare our work against NILM.

Second, we compare our approach against standard tensor factorisation (STF) introduced earlier in the paper. It should be noted that in TTF we had the following dimensions of $\mathbf{H}$, $\mathbf{A}$, and $\mathbf{S}$ respectively: $M \times h$, $h \times N \times s$, and $s \times T$. However, in STF, the dimensions of $\mathbf{H}$, $\mathbf{A}$, $\mathbf{S}$ are: $M \times r$, $N \times r$, and $T \times r$, where r is the rank of the energy breakdown tensor.

Evaluation Metric

Our evaluation metric is based on the prior work (Batra et al. 2017). It indicates how close the predicted energy breakdown is to ground-truth energy breakdown. We calculate the percentage of energy correctly assigned (PEC), where, PEC for the home, appliance, month ($\langle h, n, m \rangle$) triplet is given by:

$$PEC(h, n, m) = \frac{|E(h, n, m) - \hat{E}(h, n, m)|}{\hat{E}(h, \text{aggregate}, m)} \times 100\% \quad (5)$$

where $E(h, n, m)$ and $\hat{E}(h, n, m)$ denote the predicted and ground-truth usage by appliance n in home h in month m and $\hat{E}(h, \text{aggregate}, m)$ denotes the ground truth aggregate home energy usage for home h in month m . The RMS error in the percentage of energy correctly assigned (PEC), for an appliance n is given as the RMS of $PEC(h, n, m)$ across different months and homes,

$$RMS\ PEC(n) = \sqrt{\frac{\sum_h \sum_m PEC(h, n, m)^2}{M \times T}} \quad (6)$$

where M and T indicate the number of homes and months respectively. Lower RMS error in percentage of energy correctly assigned (PEC) means better prediction.

While this metric would allow us to evaluate the performance on a per-appliance basis; we introduce a single metric which is a weighted sum of PECs for different appliances. The weighting is done by the proportion of energy contributed by each appliance in a particular region. Such scheme would require that appliances which contribute more energy to the aggregate are more accurately estimated,

$$Weighted\ PEC = \frac{\sum_n (Frac(n) \times RMS\ PEC(n))}{\sum_n Frac(n)} \quad (7)$$

Experimental setup

Our two main experiments involve investigating the performance of transfer learning from source region Austin to target region San Diego homes, and vice-versa. Both these cities have a very different appliance energy breakdown as shown in Figure 1. When transferring $\mathbf{A}$ from Austin to San Diego, we use all the 534 homes from Austin to learn $\mathbf{A}$ and testing on all 39 homes from San Diego. However, when we perform the reverse experiment – transferring from San




Figure 1: Energy breakdown across San Diego and Austin across the 12 months in year 2014

Diego to Austin, we only use 40 test homes from Austin. The rationale is that if we have 500+ homes in a target region, and only around 40 from the source region, it would defeat the purpose of transfer learning. The 40 homes from Austin used for testing were randomly selected, and we repeated this procedure ten times to avoid bias in data sampling. As done in prior literature (Batra et al. 2017), we perform our analysis on six appliances – heating, ventilation and air-conditioning (HVAC), fridge, washing machine (WM), microwave (MW), dish washer (DW) and oven. All of these appliances have data from a significant number of homes across both San Diego and Austin. Further, these appliances also represent a wide variety of appliances. For example-season dependent (HVAC) v/s season independent (DW), background (fridge) v/s interactive (WM), etc.

Both TTF and STF involve solving $\text{Min}||E - HAS||$. STF can be solved via implementations of Canonical Polyadic Decomposition (CPD or PARAFAC) (Kolda and Bader 2009; Bro 1997) or via any of the more recent implementations (Kuleshov, Chaganty, and Liang 2015). However, we found that our implementation, which we used to solve TTF using Autograd for gradient computation, and Adagrad as the optimisation algorithm, provides the best performance over our datasets, and thus we used our projected gradient based method for solving both STF and TTF.

We use nested-cross validation across all our baselines and our approach. For the outer loop (looping across homes), we use 10-fold cross validation. The central point of investigation is: how much adaptation data do we need to beat the baselines? Thus, instead of using 9 folds for training and testing on the 10th fold, we train on $x\%$ data from the 9 folds, where $x\%$ denotes the percentage of adaptation data used. We varied x in $\{6, 7, 8, 9, 10, 20, 30, \dots, 100\}$. The rationale behind starting with $x = 6\%$ and not a lower number is that we want at least a few homes in the inner loop validation set. We randomly choose the $x\%$ of adaptation homes from the 9 folds for 5 times. This would help reduce the variance in the training set. For the inner loop, we use 2-fold cross-validation. The inner loop is used for parameter/hyper-parameter fine tuning.

The set of parameters in TTF (both normal and transfer) are: number of home and season factors; and the hyper-parameters are: the learning rate and the number of iterations. The candidate set of hyper-parameters for STF is the




Figure 2: Our proposed <TTF, Transfer> beats MF and STF baselines for Austin to San Diego transfer.




Figure 3: Our proposed <TTF, Transfer> beats MF and STF baselines for San Diego to Austin transfer.

same as that of TTF. For the STF, there is only one parameter – the rank (r).

For the MF based baselines, we used the CVXPY (Diamond and Boyd 2016) based implementation used by the paper authors. Their implementation solved the MF problem via alternating least squares. The set of parameter for <MF, Transfer> and <MF, Normal> is the number of latent factors. The set of hyper-parameters is the number of iterations of the alternating least squares.

Finally, the $Frac(n)$ required in Eq (7) are used from Table 1.

More details about the optimal parameters and hyper-parameters can be found in the Appendix. Our entire codebase, baselines, analysis and experiments can be found on Github (link anonymised for submission).

Results and Analysis

Our main result in Figure 2 shows that our approach <TTF, Transfer> performs favourably when compared to all the other baselines on Austin to San Diego transfer. Its most significant advantage over the baselines occurs for low %




Figure 4: Home factors of Austin in normal learning v.s. San Diego learnt via transfer learning from Austin.

adaptation homes, which highlights the efficacy of our approach in regions with little instrumentation. Only around 60% adaptation data does the normal learning’s error rates approach the error rate of $\langle \text{TTF}, \text{Transfer} \rangle$. This trend is expected since no transfer is perfect, and with sufficient amount of training data available in the target domain, normal learning might do just as well. We can observe that $\langle \text{STF}, \text{Normal} \rangle$ does better than $\langle \text{TTF}, \text{Normal} \rangle$. We believe this is due to the trade-off between STF and TTF with respect to the number of parameters. STF requires fewer parameters and is less prone to over-fitting when evaluated in a single region. Both $\langle \text{MF}, \text{Transfer} \rangle$ and $\langle \text{MF}, \text{Normal} \rangle$ show poor accuracy. While $\langle \text{MF}, \text{Normal} \rangle$ shows significant improvements as we increase the % of adaptation homes, the error rates remain higher than our proposed approach. It should be mentioned that previously $\langle \text{MF}, \text{Normal} \rangle$ had only been tested using a larger dataset (Batra et al. 2017). $\langle \text{MF}, \text{Transfer} \rangle$ shows little improvement wrt an increase in the % adaptation homes. We believe this is because the learnt appliance-season factor from the source domain has a strong bias towards the seasonal behaviour of the source domain.

On similar lines, results in Figure 3 show that our approach $\langle \text{TTF}, \text{Transfer} \rangle$ performs favourably compared to all the other baselines on San Diego to Austin transfer. Unlike the earlier results from Austin to San Diego transfer, in this case, none of the baselines seem to give comparable accuracy beyond 25% adaptation homes.

Having analysed and found out that our approach performs favourably compared to the baselines, we now analyse the kind of homes which will get good and bad energy breakdown estimates in the transfer setting.

Relationship between $\mathbf{H}$ and generality of our approach

The fundamental question we wish to answer in this section is: what kind of homes will receive an accurate en-

ergy breakdown under transfer settings, and what kind will receive poor energy breakdown estimates? To answer this question, we look at the learnt $\mathbf{H}$ factors. As a case-study, we take the case when we transfer $\mathbf{A}$ factors from Austin to San Diego and try to find out homes from San Diego which receive an accurate breakdown. Since the $\mathbf{A}$ factors are learnt from Austin, we compare the home factors of San Diego homes (learnt under transfer settings) to home factors of Austin homes (learnt under normal settings). The $\mathbf{H}$ are learnt with 3 home factors ($h=3$), 2 seasonal factors ($s=2$) and 10% adaptation data from San Diego (for transfer). These parameters give the best energy breakdown performance under the given setting.

Figure 4 visualises the home factors we learnt from normal learning in Austin and transfer learning in San Diego. We use K-means to cluster them into 7 clusters and plot the two most important dimensions of $\mathbf{H}$, where importance is decided by dimensions having higher variance. Of particular interest, there are a few kinds of homes:

- There is a light green cluster on the right in Figure 4 that only contains homes San Diego, which means these homes do not have similar counterparts in Austin (in $H1$ and $H2$ dimension). We observed that these set of homes have the highest energy breakdown errors.
- There is a cluster of homes on the left bottom corner containing homes from San Diego and Austin. We observed that the San Diego homes in this cluster have the lowest energy breakdown errors.
- There is a cluster of homes (light orange) which contains only Austin homes ($0.2 < H2 < 0.4$ and $0 < H1 < 0.2$). Since this cluster does not have any San Diego homes, these set of homes are not particularly useful for San Diego from a transfer learning perspective. This raises an important question, which we plan to answer in the future: given a source and a target region, can we identify the most useful homes from the source region to learn region-independent models?

We further drilled down the homes showing poor accuracy in San Diego (light green cluster towards the right side of Figure 4). We found that these homes show a high amount of HVAC error. Interestingly, none of the homes in this cluster showed a high error for the fridge. Thus, we can conclude that the homes in this cluster are different from Austin homes from an HVAC perspective.

Intuitive understanding into $\mathbf{S}$ factors Having understood the role of $\mathbf{H}$ in energy breakdown, we now look at the $\mathbf{S}$ factors learnt from our approach. In Figure 5, we look at one of the season factors learnt from San Diego and Austin under normal learning setting. We can observe that across both regions, $\mathbf{S}$ increases in the summer months and has a low value in the winter months. We also plotted the cooling degree days (CDD) in the respective regions and found a high correlation (Pearson coefficient of appx. 0.98) between this season factor and CDD. CDD is a metric to measure the amount of cooling required in a particular region¹. Without any external supervision, our approach can

¹<http://www.degreedays.net/>




Figure 5: One of the season factors learnt in our approach corresponds to the amount of air conditioning required




Figure 6: Distribution of optimal set of parameters and hyper-parameters for transfer and normal learning for TTF

learn physically relevant parameters such as a factor encoding CDD. The HVAC energy consumption is highly correlated with CDD, and given that we can learn a season factor proportional to CDD, suggests that for HVAC we can learn a region-independent model. Our approach also learns a season factor that is appx. weather independent (i.e. constant across different months). Such a season factor can encode the seasonal energy consumption of appliances such as washing machines, which do not have a direct relationship with CDD.

Limitations and Future Work

One of the most important limitations of our current work is that it will only work well if the target domain has a similar set of appliances as the source domain. If there is an unseen appliance in the target domain, we can not estimate its energy consumption. It must be pointed out that all the baselines discussed also have the same limitation. In the future, we plan to:

- **Exploit Sparsity:** Our current approach does not have any assumption about the sparsity of learnt factors. However,

we can enforce each appliance to be only affected by a small set of season and home factors (Kolter, Batra, and Ng 2010; Elhamifar and Sastry 2015), this would potentially help us address the problem of overfitting and reduce the parameter space.

- **Energy breakdown estimation for zero adaptation homes:** Our current approach requires a small number of adaptation homes from the target region. However, we believe it is possible to estimate the energy breakdown of a target region without any adaptation homes. The basic premise is that we always have the aggregate energy readings available for target region test homes; using which we can learn the home and the season factors.
- **Leveraging static features:** External metadata has been shown to be useful for improving the accuracy in collaborative filtering domain (Gu, Zhou, and Ding 2010; Rendle et al. 2011). Previous research has also shown the utility of static household information (such as household area, number of occupants, etc.) in improving energy breakdown estimates (Batra et al. 2017). In the future, we plan to leverage similar techniques, whereby, we can encourage homes with similar household properties to have similar $\mathbf{H}$ factors (Cai et al. 2011).
- **Multi-source transfer:** Currently, our approach handles the case of single source, single target. However, we believe that we can extend our approach to learn $\mathbf{A}$ from multiple regions by considering a weighted sum of objective functions from the different source regions.

Conclusions

The energy breakdown community has been looking at ways to scale across a large number of homes. One of the major bottlenecks has been that we need data from each region of interest to be able to provide an energy breakdown in that region. We believe that ours is the first approach which does not or rather requires a tiny number of homes from a target region to produce an energy breakdown. Since our evaluation proved that our approach performs favourably compared to the state-of-the-art, we believe that our approach has the potential to help scale energy breakdown.

Appendix

Figure 6 shows the distribution of the optimal set of parameters and hyper-parameters for transfer and normal learning for TTF. There are various trends. For example, for Austin to San Diego transfer, the optimal number of iterations reduces to 100 as more adaptation data is made available from San Diego. This might indicate that since more information is available from the target domain, it may be better to start with a weaker prior (i.e. learn less from the source domain using lesser number of iterations). We can also see in San Diego normal learning that as the amount of data from source domain increases, we can fit models using more number of home factors. In the interest of space, we omit the results about parameters and hyper-parameters in other settings.

References

- Armel, K. C.; Gupta, A.; Shrimali, G.; and Albert, A. 2013. Is disaggregation the holy grail of energy efficiency? The case of electricity. *Energy Policy* 52:213–234.
- Batra, N.; Wang, H.; Singh, A.; and Whitehouse, K. 2017. Matrix factorisation for scalable energy breakdown. In *AAAI*, 4467–4473.
- Batra, N.; Singh, A.; and Whitehouse, K. 2015. If you measure it, can you improve it? exploring the value of energy disaggregation. In *Proceedings of the second ACM International Conference on Embedded Systems For Energy-Efficient Built Environments*. ACM.
- Batra, N.; Singh, A.; and Whitehouse, K. 2016. Gemello: Creating a detailed energy breakdown from just the monthly electricity bill. In *SIGKDD 2016*.
- Batra, N.; Singh, A.; and Whitehouse, K. 2017. *Systems and analytical techniques towards practical energy breakdown for homes*. Ph.D. Dissertation, IIIT-Delhi.
- Bonilla, E. V.; Chai, K. M.; and Williams, C. 2008. Multi-task gaussian process prediction. In *Advances in neural information processing systems*, 153–160.
- Bro, R. 1997. Parafac. tutorial and applications. *Chemometrics and intelligent laboratory systems* 38(2):149–171.
- Cai, D.; He, X.; Han, J.; and Huang, T. S. 2011. Graph regularized nonnegative matrix factorization for data representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33(8):1548–1560.
- Dai, W.; Yang, Q.; Xue, G.-R.; and Yu, Y. 2007. Boosting for transfer learning. In *Proceedings of the 24th international conference on Machine learning*, 193–200. ACM.
- Darby, S. 2006. The effectiveness of feedback on energy consumption. *A Review for DEFRA of the Literature on Metering, Billing and direct Displays*.
- DeBruin, S.; Ghena, B.; Kuo, Y.-S.; and Dutta, P. 2015. Powerblade: A low-profile, true-power, plug-through energy meter. In *Sensys 2015*.
- Diamond, S., and Boyd, S. 2016. Cvxpy: A python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*.
- Duchi, J.; Hazan, E.; and Singer, Y. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* 12(Jul):2121–2159.
- Elhamifar, E., and Sastry, S. 2015. Energy disaggregation via learning powerlets and sparse coding.
- Gu, Q.; Zhou, J.; and Ding, C. 2010. Collaborative filtering: Weighted nonnegative matrix factorization incorporating user and item graphs. In *Proceedings of the 2010 SIAM International Conference on Data Mining*, 199–210. SIAM.
- Gupta, S.; Reynolds, M. S.; and Patel, S. N. 2010. Electrisense: single-point sensing using emi for electrical event detection and classification in the home. In *Ubicomp*.
- Harshman, R. A. 1970. Foundations of the parafac procedure: models and conditions for an “explanatory” multimodal factor analysis.
- Hart, G. W. 1992. Nonintrusive appliance load monitoring. *Proceedings of the IEEE* 80(12):1870–1891.
- Jiang, X.; Dawson-Haggerty, S.; Dutta, P.; and Culler, D. 2009. Design and implementation of a high-fidelity ac metering network. In *Information Processing in Sensor Networks, 2009. IPSN 2009. International Conference on*, 253–264. IEEE.
- Kim, Y.; Schmid, T.; Charbiwala, Z. M.; and Srivastava, M. B. 2009. Viridiscopes: design and implementation of a fine grained power monitoring system for homes. In *Proceedings of the 11th international conference on Ubiquitous computing*, 245–254. ACM.
- Kolda, T. G., and Bader, B. W. 2009. Tensor decompositions and applications. *SIAM review* 51(3):455–500.
- Kolter, J. Z., and Jaakkola, T. 2012. Approximate Inference in Additive Factorial HMMs with Application to Energy Disaggregation. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*.
- Kolter, J. Z.; Batra, S.; and Ng, A. Y. 2010. Energy Disaggregation via Discriminative Sparse Coding. In *NIPS 2010*.
- Kuleshov, V.; Chaganty, A.; and Liang, P. 2015. Tensor factorization via matrix factorization. In *Artificial Intelligence and Statistics*, 507–516.
- Lin, C.-J. 2007. Projected gradient methods for nonnegative matrix factorization. *Neural computation* 19(10):2756–2779.
- Maclaurin, D.; Duvenaud, D.; and Adams, R. P. 2015. Autograd: Effortless gradients in numpy. In *ICML 2015 AutoML Workshop*.
- Pan, S. J., and Yang, Q. 2010. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering* 22(10):1345–1359.
- Parson, O.; Ghosh, S.; Weal, M.; and Rogers, A. 2012. Non-intrusive load monitoring using prior models of general appliance types. In *AAAI 2012*.
- Parson, O.; Fisher, G.; Hersey, A.; Batra, N.; Kelly, J.; Singh, A.; Knottenbelt, W.; and Rogers, A. 2015. Dataport and nilmtk: A building data set designed for non-intrusive load monitoring. In *GlobalSIP 2015*. IEEE.
- Pérez-Lombard, L.; Ortiz, J.; and Pout, C. 2008. A review on buildings energy consumption information. *Energy and buildings* 40(3):394–398.
- Raina, R.; Battle, A.; Lee, H.; Packer, B.; and Ng, A. Y. 2007. Self-taught learning: transfer learning from unlabeled data. In *Proceedings of the 24th international conference on Machine learning*, 759–766. ACM.
- Rendle, S.; Gantner, Z.; Freudenthaler, C.; and Schmidt-Thieme, L. 2011. Fast context-aware recommendations with factorization machines. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*, 635–644. ACM.
- Shao, H.; Marwah, M.; and Ramakrishnan, N. 2013. A temporal motif mining approach to unsupervised energy disaggregation: Applications to residential and commercial buildings. In *AAAI*.
- Tucker, L. R. 1966. Some mathematical notes on three-mode factor analysis. *Psychometrika* 31(3):279–311.
- Zoha, A.; Gluhak, A.; Imran, M. A.; and Rajasegarar, S. 2012. Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey. *Sensors* 12(12):16838–16866.