

Active Collaborative Sensing for Energy Breakdown

Yiling Jia
Department of Computer Science
University of Virginia
Charlottesville, VA, USA
yj9xs@virginia.edu

Nipun Batra
Department of Computer Science
IIT Gandhinagar
Gandhinagar, India
nipun.batra@iitgn.ac.in

Hongning Wang
Kamin Whitehouse
Department of Computer Science
University of Virginia
Charlottesville, VA, USA
{hw5x,whitehouse}@virginia.edu

ABSTRACT

Residential homes constitute roughly one-fourth of the total energy usage worldwide. Providing appliance-level energy breakdown has been shown to induce positive behavioral changes that can reduce energy consumption by 15%. Existing approaches for energy breakdown either require hardware installation in every target home or demand a large set of energy sensor data available for model training. However, very few homes in the world have installed sub-meters (sensors measuring individual appliance energy); and the cost of retrofitting a home with extensive sub-metering eats into the funds available for energy saving retrofits. As a result, strategically deploying sensing hardware to maximize the reconstruction accuracy of sub-metered readings in non-instrumented homes while minimizing deployment costs becomes necessary and promising.

In this work, we develop an active learning solution based on low-rank tensor completion for energy breakdown. We propose to actively deploy energy sensors to appliances from selected homes, with a goal to improve the prediction accuracy of the completed tensor with minimum sensor deployment cost. We empirically evaluate our approach on the largest public energy dataset collected in Austin, Texas, USA, from 2013 to 2017. The results show that our approach gives better performance with fixed number of sensors installed, when compared to the state-of-the-art, which is also proven by our theoretical analysis.

CCS CONCEPTS

• **Computing methodologies** → **Active learning settings**; • **Human-centered computing** → *Ubiquitous computing*.

KEYWORDS

active learning, tensor completion, energy breakdown

ACM Reference Format:

Yiling Jia, Nipun Batra, Hongning Wang, and Kamin Whitehouse. 2019. Active Collaborative Sensing for Energy Breakdown. In *The 28th ACM International Conference on Information and Knowledge Management (CIKM '19)*, November 3–7, 2019, Beijing, China. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3357384.3357929>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM '19, November 3–7, 2019, Beijing, China

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6976-3/19/11...\$15.00

<https://doi.org/10.1145/3357384.3357929>

1 INTRODUCTION

Residential homes are one of the largest energy consumers worldwide, constituting roughly one-fourth of the total energy usage [31]. Part of this energy could be saved by providing an energy breakdown, i.e., per-appliance energy consumption summary. Studies have shown that energy breakdown enables informed decision making by different actors in the home's energy ecosystem [1]. For example, studies [1, 22] report energy feedback causes behavioral changes that can reduce energy consumption by 15%. It also helps power utility companies and policymakers to improve load forecasting [1], detect broken or mis-configured equipment [20], and target the most inefficient homes for energy efficiency programs.

Various energy breakdown techniques have been proposed in the past, such as direct sensing systems [14, 19] and non-intrusive load monitoring (NILM) [16, 17, 24, 25]. Different from those models which require hardware to be installed in every home, recently, collaborative sensing [5–7] has attracted increasing attention due to its low cost and high scalability. Collaborative sensing, which aims at reconstructing the appliance-level energy data of non-instrumented homes based on data collected from other homes, only requires the easily available information of non-instrumented homes, such as the monthly energy bills, square footage, and number of occupants. The basic premise is that, while every home is unique, the common design and construction patterns among homes create a shared and repeating structure, which gives rise to a sparse set of factors contributing to energy variation across homes. A typical approach is to factorize energy readings into a low-dimensional space, and predict energy consumption in a non-instrumented home with this low-dimensional model based on the high-fidelity data collected in other homes. It is worth mentioning that collaborative sensing algorithms only perform at a low temporal resolution, such as monthly data, unlike NILM-type solutions that can produce a high frequency appliance energy time-series. However, previous studies have shown sustained savings even at a monthly resolution [15, 23], which support the value of low frequency energy breakdown.

While collaborative sensing alleviates the scalability issue of NILM-type methods by removing the requirement of per-home instrumentation, these solutions have assumed the existence of relevant training data, i.e., appliance-level energy readings, from some fully instrumented homes. But, in reality, very few homes in the world have been instrumented with sub-meters (appliance level energy meters). As a result, sensor deployment is still inevitable to apply such methods. In this work, to further improve the scalability of collaborative sensing, we seek to answer the question: *can we minimize the deployment cost by selectively deploying sensing hardware to a subset of homes and appliances while maximizing the*




Figure 1: Active sensor deployment for energy breakdown. (1) We treat aggregate readings of each home as a special appliance, and they are always available; (2) At the end of January, if $\langle \text{home } x_1, \text{appliance } y_1 \rangle$ is selected, appliance readings $e_{x_1, y_1, t}$ become available for $t \geq 2$.

reconstruction accuracy of sub-metered readings in non-instrumented homes? We name this new research problem as active sensor deployment for energy breakdown.

Active sensor deployment for energy breakdown differs from classical active learning problems [13, 33, 38] in three major aspects. First, energy readings are time-series data. New readings are constantly generated, and they are influenced by various external and internal factors, such as season [6] and occupant activities. The typically imposed assumption in active learning literature that observations are independent and identically distributed no longer holds in this situation. Second, once a sensor has been installed in a home, the monitored appliance readings from that home will become available thereafter. This directly introduces the explore/exploit dilemma in active sensor deployment, because one has to balance the choice of instrumentation that focuses on the current bottleneck of reconstruction accuracy, and that improves model accuracy for future predictions. For example, in spring, furnace might consume most energy per-household for heating, and therefore more instrumentation on furnace is required to obtain a more accurate model. But in summer, air conditioning system will become the main source of energy consumption. If one has not instrumented any air conditioning system before, the reconstruction accuracy on it will be poor. It is necessary to plan the instrumentation ahead of time, so as to obtain a high accuracy model before the consumption peaks. Third, an instrumentation choice concerns two different types of entities, i.e., homes and appliances, which are not independent. This adds another dimension into the explore/exploit dilemma for sequential decision making: which $\langle \text{home}, \text{appliance} \rangle$ pair to instrument next for maximizing future reconstruction accuracy.

In this work, we follow factorization based collaborative sensing [6, 7], and propose to perform active sensor deployment via active tensor completion. First of all, we view energy readings as a three-way tensor as illustrated in Figure 1: with homes, appliances and time as three separate dimensions. In this tensor, a cell is filled with an observation if the corresponding appliance has been monitored in this home before; and predictions are made in cells with missing values. Although the size of this tensor is large and

increasing, due to the continuously observed energy readings, we believe the tensor is low rank. Learning a low rank representation for homes and appliances enables us to predict energy use in those non-instrumented homes. Evidence of this modeling assumption has been gathered in various existing studies [6, 8]. As a result, the active sensor deployment problem can be naturally formalized as an active tensor completion problem: deciding which cells in the tensor to query so that the reconstruction accuracy can be maximized. Specifically, at the end of each month, we query the home and appliance pairs that have the highest uncertainty in the current tensor reconstruction, which we prove to reduce reconstruction uncertainty mostly rapidly. And to project a model's prediction uncertainty of future readings in a longer term, we incorporate external seasonal information into model estimation. This helps the model react to future season changes earlier. We name our solution as Active Collaborative Sensing, or ActSense in short.

We rigorously prove that with high probability the developed algorithm achieves a considerable reduction in accumulated estimation error. In other words, it requires less sensor deployment to obtain the same level of reconstruction accuracy compared to any other sensor installation. We evaluate ActSense on a publicly available dataset called Dataport [29], which is the largest public energy dataset collected in U.S. With a fixed budget of sensors, our model achieves better energy breakdown performance compared to three baseline approaches. Besides, extensive analysis of the experiments shows that integrating the temporal seasonal information can help to foresee the energy usage trends and prepare the sensor installation in advance.

2 RELATED WORK

Since George Hart's seminal work on non-intrusive load monitoring (NILM) in the early 1990s [17], the research community has proposed several solutions to scale up energy breakdown. The basic idea of NILM algorithms is to perform source separation on the power signal measured at a single point (home mains) [1, 42]. Majority of these algorithms worked on time-series data obtained from a smart meter, collected at rates from 10s of kHz to a reading once every hour [24, 25, 30, 34]. But NILM type algorithms still require instrumentation across each home. Recently, there is a line of research collectively referred to as collaborative sensing that focuses on energy breakdown without any hardware installation in a test home [5, 6, 8]. Their key idea is that "similar homes would have a similar per-appliance energy consumption". These approaches estimate the energy breakdown of a home by finding a similar home (based on monthly energy use) that already has an energy breakdown available. These approaches were shown to be scalable and accurate compared to the state-of-the-art NILM approaches. However, such algorithms require available energy breakdown data for model learning; they cannot apply when this data is missing.

The active sensor deployment problem has been studied in robotics for several decades. But the main focus there is on active sensing in a fixed region by utilizing the spatial information. For example, in [12], the vision sensor was purposefully configured and placed at several positions to observe a target. Ksenia et al. [35] addressed the object search problem in surveillance by optimizing the probability of finding the target given a fixed cost limit. In [4, 11, 32],

the authors proposed the approaches with next best view planning methods to improve object modeling and recognition, while several studies have been explored on active sensing based on information entropy and some rule-based methods [26, 28]. But in active energy breakdown problem that we study in this paper, the energy data is a time series generated continuously from different homes and appliances. Temporal pattern in energy consumption across different homes and appliances is the key for optimizing sensor deployment. To the best of our knowledge, no existing work addresses the active sensing problem in energy breakdown.

On a related direction of research, active matrix completion has gained increasing attention in the past decade, due to the wide adoption of matrix completion based collaborative filtering algorithms. The essence of active matrix completion is to assess statistical uncertainty of each unobserved cell in the matrix, and propose to query the most informative ones to improve the factorization model's prediction accuracy. Variational Bayesian Inference (VBI) [9] is one of the most popularly used methods for uncertainty assessment. There exist several solutions leveraging VBI to estimate model uncertainty and using mutual information [36], prediction variance [37] and information gain [18] to query the observations actively. Shayok et al. utilized graphical lasso for model estimation and queried instances with the highest uncertainty [10]. Some recent developments solve online matrix completion with multi-armed bandit algorithms, a reference solution for explore-exploit trade-off [2, 3, 27]. Model uncertainty is assessed by posterior sampling [21] or estimation of prediction's confidence interval [40]. However, in the problem of energy breakdown, it is insufficient to view the energy data as a matrix, as the data is continuously generated over time. Observations collected from different time periods now become depend on each other. An optimal query strategy should take this temporal dependency into consideration to balance the uncertainty from different homes and appliances over time. This, to the best of our knowledge, has not been carefully studied to date.

3 METHODOLOGY

There are two essential research questions in active sensor deployment for energy breakdown: (1) how to accurately perform energy breakdown on a given set of homes with sub-metering data; and (2) how to select <home, appliance> pairs for additional instrumentation to improve future energy breakdown accuracy. In this section, we first provide a detailed problem definition, then elaborate the design of our proposed Active Collaborative Sensing solution, and provide theoretical analysis of its convergence in the end.

3.1 Problem Statement

In active sensor deployment, we aim to maximize the energy breakdown accuracy while minimizing the sensor deployment cost. We formally define our energy reading data as a three-way tensor, $\mathbf{E}_{M \times N \times T}$, where the elements contain energy readings from M homes and N appliances for up to T months. Detailed reading is available in a cell when the corresponding appliance is monitored in that home beforehand; otherwise the reading is unknown. There are some special structural properties of this tensor, which make the active sensor deployment problem especially challenging.

- **Time-series data.** In energy breakdown scenario, energy data is continuously generated and collected after every sampling cycle, e.g., monthly in Figure 1. At the end of each month, new <home, appliance> pairs can be selected according to predefined selection strategy and sensor deployment budget.

- **Sensor installation.** In sensor deployment, once a sensor is installed, the appliance readings will always be available in the future. For example, if we install K sensors in January and K sensors in February, we will have K extra readings at the end of January and $2 \times K$ extra readings at the end of February.

- **Aggregate readings.** We view the household aggregate energy as a "special" appliance in the energy tensor, which is always available in the form of monthly energy bill.

The active sensor deployment for energy breakdown can be formalized as an active tensor completion problem. More specifically, the learning procedure is: at month $t \in [1, T - 1]$, based on the observed energy tensor $\mathbf{E}_\Omega \in \mathbf{R}^{M \times N \times t}$ where Ω represents the set of indices of observed elements, and subject to budget constraint, we want to install sensors to a set of <home, appliance> pairs which will bring the largest reconstruction error reduction of the energy tensor in the future months.

3.2 Energy Tensor Completion

The basic factorization model we adopt here is the Canonical Polyadic (CP) decomposition, which is also known as rank decomposition. Our core intuition is that the common design and construction patterns in residential homes create a shared and repeating structure. And the heterogeneity across homes can be captured by a low dimension representation of individual homes [6]. Examples of such low-dimensional representations could be home insulation or the number of occupants. Similarly, the energy consumption dependence of different appliances with respect to home design and weather pattern can also be encoded using a low-dimensional representations.

Therefore, we can decompose the energy tensor $\mathbf{E}_{M \times N \times T}$ into three factors: home factor $\mathbf{H} \in \mathbf{R}_+^{M \times r}$, appliance factor $\mathbf{A} \in \mathbf{R}_+^{N \times r}$ and season factor $\mathbf{S} \in \mathbf{R}_+^{T \times r}$, where $\mathbf{H}$, $\mathbf{A}$ and $\mathbf{S}$ are independent non-negative matrices and r is the rank of the energy tensor. With such a decomposition design, the i th row of $\mathbf{H}$, denoted as $\mathbf{h}_i$, is the latent factor of the i th home; the j th row of $\mathbf{A}$, denoted as $\mathbf{a}_j$, is the latent factor for j th appliance; and the k th row of $\mathbf{S}$, denoted as $\mathbf{s}_k$, is the latent factor for k th month. To reflect the fact that energy readings should be finite in all homes and appliances, we further assume that the L2 norm of these latent matrices are upper bounded: $\forall i, \|\mathbf{h}_i\|_2 \leq P$; $\forall j, \|\mathbf{a}_j\|_2 \leq Q$; and $\forall k, \|\mathbf{s}_k\|_2 \leq R$. The energy consumption of appliance j in home i at month k can be estimated by $\mathbb{E}[\mathbf{e}_{ijk}] = \langle \mathbf{h}_i, \mathbf{a}_j, \mathbf{s}_k \rangle$, where $\langle \cdot, \cdot, \cdot \rangle$ represents the triple product. In each month t , in order to estimate the latent factors, regularized quadratic loss over the observations in the energy tensor $\mathbf{E}_{M \times N \times t}$ is employed. The objective function for latent factor estimation at month t is,

$$F = \sum_{(i,j,k) \in \Omega_t} (\langle \mathbf{h}_i^t, \mathbf{a}_j^t, \mathbf{s}_k^t \rangle - \mathbf{e}_{ijk})^2 \quad (1)$$

$$+ \lambda_1 \sum_{i=1}^M \|\mathbf{h}_i^t\|_2 + \lambda_2 \sum_{j=1}^N \|\mathbf{a}_j^t\|_2 + \lambda_3 \sum_{k=1}^t \|\mathbf{s}_k^t\|_2$$

where Ω_t represents the set of indices of observed energy readings till month t , and λ_1, λ_2 and λ_3 are the coefficients for L2 regularization. The notion of $(\mathbf{h}_i^t, \mathbf{a}_j^t, \mathbf{s}_k^t)$ denotes the latest estimate of latent factors at month t . The inclusion of L2 regularization terms is critical to our solution in two folds. First, it makes the sub-problems in coordinate decent based optimization well-posed, so that we have closed form solutions for the latent factors at each round. Second, it helps to remove the scaling indeterminacy among the estimates of those latent factors.

We estimate the latent factors with alternative least square (ALS). Specifically, the closed-form estimation of $(\mathbf{h}_i, \mathbf{a}_j, \mathbf{s}_k)$ with respect to Eq (1) at month t can be achieved by $\hat{\mathbf{h}}_i^t = (\mathbf{A}_i^t)^{-1}\mathbf{b}_i^t$, $\hat{\mathbf{a}}_j^t = (\mathbf{C}_j^t)^{-1}\mathbf{d}_j^t$ and $\hat{\mathbf{s}}_k^t = (\mathbf{E}_k^t)^{-1}\mathbf{f}_k^t$, in which,

$$\mathbf{A}_i^t = \lambda_1 \mathbb{I}_1 + \sum_{(i,j,k) \in \Omega_t} (\mathbf{a}_j^t \circ \mathbf{s}_k^t)(\mathbf{a}_j^t \circ \mathbf{s}_k^t)^\top, \mathbf{b}_i^t = \sum_{(i,j,k) \in \Omega_t} \mathbf{e}_{ijk}(\mathbf{a}_j^t \circ \mathbf{s}_k^t) \quad (2)$$

$$\mathbf{C}_j^t = \lambda_2 \mathbb{I}_2 + \sum_{(i,j,k) \in \Omega_t} (\mathbf{h}_i^t \circ \mathbf{s}_k^t)(\mathbf{h}_i^t \circ \mathbf{s}_k^t)^\top, \mathbf{d}_j^t = \sum_{(i,j,k) \in \Omega_t} \mathbf{e}_{ijk}(\mathbf{h}_i^t \circ \mathbf{s}_k^t) \quad (3)$$

$$\mathbf{E}_k^t = \lambda_3 \mathbb{I}_3 + \sum_{(i,j,k) \in \Omega_t} (\mathbf{h}_i^t \circ \mathbf{a}_j^t)(\mathbf{h}_i^t \circ \mathbf{a}_j^t)^\top, \mathbf{f}_k^t = \sum_{(i,j,k) \in \Omega_t} \mathbf{e}_{ijk}(\mathbf{h}_i^t \circ \mathbf{a}_j^t) \quad (4)$$

where $\mathbb{I}_1, \mathbb{I}_2$ and $\mathbb{I}_3$ are identity matrices with dimensions of $r \times r$. The operation $(\circ)$ represents the element-wise product of two vectors, and it is easy to verify that $\langle \mathbf{h}_i, \mathbf{a}_j, \mathbf{s}_k \rangle = (\mathbf{h}_i \circ \mathbf{a}_j)^\top \mathbf{s}_k = (\mathbf{h}_i \circ \mathbf{s}_k)^\top \mathbf{a}_j = (\mathbf{a}_j \circ \mathbf{s}_k)^\top \mathbf{h}_i$. Proper projection is needed to ensure the estimated factors are non-negative and their norms are in the required range. The estimated factors predict the unseen values in the energy tensor at month t by $\hat{\mathbf{e}}_{ijk} = \langle \hat{\mathbf{h}}_i^t, \hat{\mathbf{a}}_j^t, \hat{\mathbf{s}}_k^t \rangle$.

3.3 Uncertainty Quantification

With the learnt latent factors, household energy breakdown can be readily provided by the aforementioned factorization-based solution. However, without sufficient appliance-level energy readings, the estimated latent factors are subject to various sources of uncertainty, e.g., variance in energy use or errors in sensor readings, which directly lead to volatility in appliance-level energy prediction. Thus, to improve the estimation quality of future energy consumption and obtain strong performance guarantee, strategically querying observations from the non-instrumented home and appliance pairs is of paramount importance. In ActSense, we propose to select the unobserved pairs with the highest estimation uncertainty so as to bring the largest reconstruction error reduction.

The first step in active selection is to quantify the estimation uncertainty in the energy tensor completion. The main reason for the estimation uncertainty is the existence of potential noise in the observed energy readings, e.g., imperfect sensor hardware, unexpected energy use (such as energy consumed in wire transition), and etc. Passing through the factorization procedure, such noise leads to uncertainty in latent factor estimation, e.g., $\|\hat{\mathbf{h}}_i^t - \mathbf{h}_i^*\| \neq 0$, where $\mathbf{h}_i^*$ is the ground-truth latent home factor, and $\hat{\mathbf{h}}_i^t$ is the home factor estimated with the observed noisy energy readings $\mathbf{e}_{ijk}^{obs}$ at month t . The discrepancy between the estimated and ground-truth latent factors directly contributes to the uncertainty in appliance-level energy consumption estimation. To model uncertainty, we assume the noise follows a zero-mean Gaussian distribution: $\mathbf{e}_{ijk}^{obs} =$

$\mathbf{e}_{ijk}^* + \eta_{ijk} = \langle \mathbf{h}_i^*, \mathbf{a}_j^*, \mathbf{s}_k^* \rangle + \eta_{ijk}$, where $\eta_{ijk} \sim \mathcal{N}(0, \sigma_{ijk}^2)$ is the noise term, $\mathbf{e}_{ijk}^{obs}$ is the observed energy reading, $\mathbf{e}_{ijk}^*$ is the noise-free energy consumption, and $\mathbf{h}_i^*, \mathbf{a}_j^*$, and $\mathbf{s}_k^*$ are the ground-truth latent factors. Under the context of tensor completion for energy breakdown, at month t , the uncertainty of appliance-level energy estimation comes from the estimation error of latent home factors, i.e., $\|\hat{\mathbf{h}}_i^t - \mathbf{h}_i^*\|$, appliance factors, i.e., $\|\hat{\mathbf{a}}_j^t - \mathbf{a}_j^*\|$, and season factors, i.e., $\|\hat{\mathbf{s}}_k^t - \mathbf{s}_k^*\|$, caused by the noise. With the closed-form solution in our coordinate descent estimation, at each time t , the confidence interval of estimated $\hat{\mathbf{h}}_i^t$, $\hat{\mathbf{a}}_j^t$, and $\hat{\mathbf{s}}_k^t$ can be analytically computed by the following lemma.

LEMMA 3.1. *With proper initialization of the coordinate descent estimation, the Hessian matrix is positive definite at the optimizer $\mathbf{h}_i^*$, $\mathbf{a}_j^*$, and $\mathbf{s}_k^*$. Thus, for any $\epsilon_1 > 0$, $\epsilon_2 > 0$, $\epsilon_3 > 0$, and $\delta \in (0, 1)$, with probability at least $1 - \delta$, the estimated error of latent factors satisfies,*

$$\begin{aligned} \|\hat{\mathbf{h}}_i^t - \mathbf{h}_i^*\|_{\mathbf{A}_i^t} &\leq \sqrt{r \ln \frac{\lambda_1 r + |\Omega_t| Q^2 R^2}{\lambda_1 \cdot r \cdot \delta}} + \sqrt{\lambda_1} P + \frac{2PQ^2 R^2}{\sqrt{\lambda_1}} (G_2 + G_3) \\ \|\hat{\mathbf{a}}_j^t - \mathbf{a}_j^*\|_{\mathbf{C}_j^t} &\leq \sqrt{r \ln \frac{\lambda_2 r + |\Omega_t| P^2 R^2}{\lambda_2 \cdot r \cdot \delta}} + \sqrt{\lambda_2} Q + \frac{2P^2 Q R^2}{\sqrt{\lambda_1}} (G_1 + G_3) \\ \|\hat{\mathbf{s}}_k^t - \mathbf{s}_k^*\|_{\mathbf{E}_k^t} &\leq \sqrt{r \ln \frac{\lambda_3 r + |\Omega_t| P^2 Q^2}{\lambda_3 \cdot r \cdot \delta}} + \sqrt{\lambda_3} R + \frac{2P^2 Q^2 R}{\sqrt{\lambda_3}} (G_1 + G_2) \end{aligned}$$

where,

$$\begin{aligned} G_1 &= \frac{f_1(1 - f_1^{|\Omega_t|})}{1 - f_1}, G_2 = \frac{f_2(1 - f_2^{|\Omega_t|})}{1 - f_2}, G_3 = \frac{f_3(1 - f_3^{|\Omega_t|})}{1 - f_3} \\ f_1 &= q_1 + \epsilon_1, \quad f_2 = q_2 + \epsilon_2, \quad f_3 = q_3 + \epsilon_3 \end{aligned}$$

and $q_1 \in (0, 1)$, $q_2 \in (0, 1)$, $q_3 \in (0, 1)$, $|\Omega|$ is the cardinality of set Ω .

The detailed proof of Lemma 3.1 can be found in the provided supplementary material¹. With Lemma 3.1, we obtain a tight construction for the estimation uncertainty of the latent factors, which can be easily transformed to the uncertainty of energy tensor estimation $\hat{\mathbf{e}}_{ijk}$. As described in Section 3.1, in active energy tensor completion, the $\langle \text{home}, \text{appliance} \rangle$ pairs are selected at the end of each month. Under such a setting, we do not have a choice of when to install the sensors. Therefore, we define α_x^t and α_y^t as the upper bound of $\|\hat{\mathbf{h}}_i^t - \mathbf{h}_i^*\|_{\mathbf{A}_i^t}$ and $\|\hat{\mathbf{a}}_j^t - \mathbf{a}_j^*\|_{\mathbf{C}_j^t}$, and propose the following $\langle \text{home}, \text{appliance} \rangle$ pair selection strategy at month t :

$$(x, y) = \arg \max_{x \in [1, M], y \in [1, N]} \left(\alpha_x^t \sqrt{(\mathbf{a}_y^t \circ \hat{\mathbf{s}}_t^t)^\top \mathbf{A}_x^t{}^{-1} (\mathbf{a}_y^t \circ \hat{\mathbf{s}}_t^t)} + \alpha_y^t \sqrt{(\mathbf{h}_x^t \circ \hat{\mathbf{s}}_t^t)^\top \mathbf{C}_y^t{}^{-1} (\mathbf{h}_x^t \circ \hat{\mathbf{s}}_t^t)} \right) \quad (5)$$

The detailed derivation of this uncertainty estimation is provided in the supplementary material. In Eq (5), the two terms on the right-hand side are related to the estimation uncertainty of latent factor $\hat{\mathbf{h}}_i$ and $\hat{\mathbf{a}}_j$ at current month t . We choose to select the pairs with the highest estimation uncertainty of the associated latent factors, which leads to the best reduction in the model's overall prediction uncertainty. We prove this conclusion in Section 3.5.

¹<http://arxiv.org/abs/1909.00525>

3.4 Active Collaborative Sensing

The aforementioned uncertainty estimation is constructed by the estimated season factor s_t^t of month t ; hence, it only measures the model's prediction uncertainty at that particular month. But it does not consider the specific challenges we face in the context of energy breakdown (as discussed in Section 3.1). Basically, there exists a time lag in sensor deployment: all the decisions we make are based on the current knowledge of energy consumption, but the target of instrumentation is to get a better energy breakdown accuracy for future predictions. We can only verify our decisions at least one month later. The active selection will be more effective if we could foresee the changes in future and prepare the sensor installation accordingly. However, with limited observations, it is hard to extrapolate the usage pattern of different appliances across homes nor to make a reasonable prediction of future use, not to mention the uncertainty estimation of this future prediction.

Motivated by the analysis in previous literature [6], within one geo-region, the season factors learned from energy data are highly correlated with the region's seasonal pattern. Previous work also showed that the season patterns are similar and repeating across years. Hence, with the aggregate readings collected from the monthly bills of the past year, we can obtain a rough estimation of the season factors of the current year. We then inject historically learned season factors into our latent factor estimation by changing Eq (1) to,

$$F = \sum_{(i,j,k) \in \Omega_t} (\langle \mathbf{h}_i^t, \mathbf{a}_j^t, \mathbf{s}_k^t \rangle - \mathbf{e}_{ijk})^2 + \lambda_1 \sum_{i=1}^M \|\mathbf{h}_i^t\|_2 \quad (6)$$

$$+ \lambda_2 \sum_{j=1}^N \|\mathbf{a}_j^t\|_2 + \lambda_3 \sum_{k=1}^t \|s_k^t - s_k^{prev}\|_2$$

where s^{prev} is the season factor learned from past years. As s^{prev} can be estimated from aggregate readings alone, we even do not require any instrumented homes from past years for this estimation. This new regularization term assumes that the season across years is similar, and it changes smoothly between two adjacent years. As a result, when the observations from future months are not yet available, we can directly use the previously estimated season factors as an estimate for those future months. With these estimated season factors, we can extend Eq (5) to quantify the model's prediction uncertainty for future months.

On the other hand, because at the end of each month we update the latent factor estimations based on the latest available energy readings, we would have an updated uncertainty estimation on the $\langle \text{home}, \text{appliance} \rangle$ pairs in the past months as well (by using the updated season factors for those months in Eq (5)), if we have not instrumented them yet. It provides us information about where our model is still uncertain in historical observations. It is therefore necessary to combine the uncertainty estimations about past, current, and future energy use to guide active sensor deployment. We use a time decay kernel to integrate those uncertainty estimations, and describe the procedure in Algorithm 1.

As shown in Algorithm 1, the final uncertainty of a $\langle \text{home}, \text{appliance} \rangle$ pair is a linear combination of uncertainties across a fixed period of time, i.e., 12 months. The $\mathbf{Weight}(t, t')$ function controls how much a particular month's uncertainty estimation

Algorithm 1: Uncertainty(x, y, t)

Input : $\mathbf{h}_x^t, \mathbf{a}_y^t, \mathbf{s}_{1:t}^t, \mathbf{A}_x^t, \mathbf{C}_y^t, \mathbf{S}^{prev}, \alpha_x, \alpha_y$
Output : $U(x, y, t)$

```

1  $U(x, y, t) = 0$ 
2 for  $t' = 1$  to  $T$  do
3   if  $t' \leq t$  then  $\tilde{s} \leftarrow s_{t'}^t$ , else  $\tilde{s} \leftarrow s_{t'}^{prev}$ 
4    $U(x, y, t) \leftarrow$ 
      $U(x, y, t) + \mathbf{Weight}(t', t) \cdot \left( \alpha_x \sqrt{(\mathbf{a}_y^t \circ \tilde{s})^\top (\mathbf{A}_x^t)^{-1} (\mathbf{a}_y^t \circ \tilde{s})} + \right.$ 
      $\left. \alpha_y \sqrt{(\mathbf{h}_x^t \circ \tilde{s})^\top (\mathbf{C}_y^t)^{-1} (\mathbf{h}_x^t \circ \tilde{s})} \right)$ 
5 end
```

contributes to the final decision. We use a triangle kernel to calculate the weight:

$$k(t', t) = \begin{cases} 1 - \frac{|t' - t|}{\sigma} & \text{if } |t' - t| \leq \sigma \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

where t is the index of current month, and t' is the index of a candidate month, and σ is the parameter to control time decay. We always give the current month itself the highest weight. The intuition behind Eq (7) is that as the season is expected to smoothly change between adjacent months, the estimation uncertainty from nearby months should be more important for the decision of sensor deployment. For example, in May, the uncertainty estimated from February should be less important than that from June.

Combining the integrated uncertainty estimation and the estimation of latent factors discussed in Section 3.2, we can effectively perform active tensor completion, and therefore address the problem of active sensor deployment for energy breakdown. We name the resulting algorithm Active Collaborative Sensing, or ActSense in short, and illustrate the detailed procedure of it in Algorithm 2. Here, we assume uniform cost for instrumenting different appliances. Thus, considering a practical setting in active sensor deployment, at the end of each month, $L \langle \text{home}, \text{appliance} \rangle$ pairs with the highest uncertainties will be selected for sensor installation, where L can be determined by the total cost of instrumentation and budget. And we leave the work of measuring the various cost of sensor installation across different appliances as our future work.

3.5 Convergence Analysis

In this section, we theoretically analyze the sample selection strategy in our proposed Active Collaborative Sensing algorithm.

Based on Lemma 3.1 and the q -linear convergence property of alternative least square [39], at month t , for any home i , appliance j , and any month k before t , i.e., $k \leq t$, the difference between our estimation and ground-truth energy reading is bounded by,

$$|\hat{\mathbf{e}}_{ijk} - \mathbf{e}_{ijk}^*| \leq \alpha_i^t \|\hat{\mathbf{a}}_j^t - \hat{\mathbf{s}}_k^t\|_{(\mathbf{A}_i^t)^{-1}} + \alpha_j^t \|\hat{\mathbf{h}}_i^t - \hat{\mathbf{s}}_k^t\|_{(\mathbf{C}_j^t)^{-1}} \quad (8)$$

$$+ 4PQR(q_2 + \epsilon_2)^{t+1} + 2PQR(q_3 + \epsilon_3)^{t+1}$$

in which α_i^t and α_j^t are the upper bound of $\|\hat{\mathbf{h}}_i^t - \mathbf{h}_i^*\|_{\mathbf{A}_i^t}$ and $\|\hat{\mathbf{a}}_j^t - \mathbf{a}_j^*\|_{\mathbf{C}_j^t}$, and they can be explicitly calculated based on Lemma 3.1. The detailed derivation is provided in our supplementary material. According to Lemma 3.1, the last two terms on the right-hand

Algorithm 2: Active Collaborative Sensing

Input : $\lambda_1, \lambda_2, \lambda_3 \in (0, \infty), L$

- 1 **foreach** $i = 1$ **to** M **do** $A_i^0 \leftarrow \lambda_1 I, b_i^0 \leftarrow 0, h_i^0 \leftarrow 0$;
- 2 **foreach** $j = 1$ **to** N **do** $C_j^0 \leftarrow \lambda_2 I, d_j^0 \leftarrow 0, a_j^0 \leftarrow 0$;
- 3 **foreach** $t = 1$ **to** T **do** $E_t^0 \leftarrow \lambda_2 I, f_t^0 \leftarrow 0, s_t^0 \leftarrow 0$;
- 4 Initialize the observation set $\Theta_0 \leftarrow \emptyset$;
- 5 **for** $t = 1$ **to** T **do**
- 6 Update Θ_t with monthly bills and sub-meter readings.;
- 7 **for** e_{ijk} from home i , appliance j at time k where $e_{ijk} \in \Theta_t$ **do**
- 8 Update A_i^t, b_i^t by Eq (2), $h_i^t \leftarrow (A_i^t)^{-1} b_i^t$;
- 9 Update C_j^t, d_j^t by Eq (3), $a_j^t \leftarrow (C_j^t)^{-1} d_j^t$;
- 10 Update E_k^t, f_k^t by Eq (4), $s_k^t \leftarrow (E_k^t)^{-1} f_k^t$;
- 11 **end**
- 12 Install sub-meters to L <home, appliance> pairs with highest **Uncertainty**(x, y, t);
- 13 **end**

side of the inequality are upper bounded, and their upper bound is independent from the procedure of sample selection. And the first two terms are calculated based on the currently estimated parameters in ActSense.

Based on Eq (8), we compare the contribution of the selection made by ActSense in reducing its uncertainty of energy breakdown with that from any other possible selections. Denote $\langle x_a, y_a \rangle$ as the indices of <home, appliance> pair selected by ActSense at the end of month t , and $\langle x_o, y_o \rangle$ as the indices of any other pairs. At month $t + 1$, because of different new observations introduced by these two selections, e.g., energy consumed by $\langle x_a, y_a \rangle$ and $\langle x_o, y_o \rangle$, we have two different energy breakdown estimations. Here, we use $(\hat{e}_{x_a, y_a, t+1}^A, \hat{e}_{x_o, y_o, t+1}^A)$ to represent the energy prediction after observing $e_{x_a, y_a, t+1}$, and $(\hat{e}_{x_a, y_a, t+1}^O, \hat{e}_{x_o, y_o, t+1}^O)$ to represent the prediction if the observation is $e_{x_o, y_o, t+1}$. Thus, for month $t + 1$, we use $E_A(t + 1)$ to represent the prediction error based on the selection made by ActSense, and $E_O(t + 1)$ for the prediction error resulted from any other selections,

$$E_A(t + 1) = |\hat{e}_{x_a, y_a, t+1}^A - e_{x_a, y_a, t+1}^*| + |\hat{e}_{x_o, y_o, t+1}^A - e_{x_o, y_o, t+1}^*|$$

$$E_O(t + 1) = |\hat{e}_{x_a, y_a, t+1}^O - e_{x_a, y_a, t+1}^*| + |\hat{e}_{x_o, y_o, t+1}^O - e_{x_o, y_o, t+1}^*|$$

Here, we only consider the estimation error contributed by these two pairs, because the errors from other estimations are bounded by the same result as shown in Eq (8).

With ActSense, at month t , $\langle x_a, y_a \rangle$ is selected as it has the highest uncertainty among all the <home, appliance> pairs. With such a condition, and the q -linear convergence property of alternative least square, it can be proved that, at month $t + 1$, for $\delta \in (0, 1)$, with probability at least $1 - \delta$, the upper bound of the estimation error generated with ActSense, i.e., $UB(E_A(t + 1))$, and the one generated with any other selections, i.e., $UB(E_O(t + 1))$, satisfy:

$$UB(E_A(t + 1)) \leq UB(E_O(t + 1)) \quad (9)$$

PROOF SKETCH. In the proof, we first need to derive the upper bound of the estimation errors and then obtain the relationship

between them. The key idea for deriving the upper bound is to obtain the estimation error at month $t + 1$ based on the learned parameters and the selection made at month t . Take the first term of $E_A(t + 1)$ as an example, $|\hat{e}_{x_a, y_a, t+1}^A - e_{x_a, y_a, t+1}^*|$ can be bounded by $\alpha_{x_a}^{t+1} \|\hat{a}_{y_a}^{t+1} \circ \hat{s}_{t+1}^{t+1}\|_{(A_{x_a}^{t+1})^{-1}} + \alpha_{y_a}^{t+1} \|\hat{h}_{x_a}^{t+1} \circ \hat{s}_{t+1}^{t+1}\|_{(C_{y_a}^{t+1})^{-1}} + C_1$, where C_1 is a constant and independent from any selections. First, we assume that the season factors change smoothly between adjacent months, e.g., $\|\hat{s}_k^t - \hat{s}_{k+1}^t\|_2 \leq \|u\|_2 = \gamma$. With such an assumption and the q -linear convergence property, we relax the upper bound with the latent factors learned at month t , e.g., $\alpha_{x_a}^{t+1} \|\hat{a}_{y_a}^{t+1} \circ \hat{s}_{t+1}^{t+1}\|_{(A_{x_a}^{t+1})^{-1}} \leq \alpha_{x_a}^{t+1} \|\hat{a}_{y_a}^t \circ \hat{s}_t^t\|_{(A_{x_a}^{t+1})^{-1}} + C_2$, where C_2 is also a constant. Since $A_{x_a}^t$ is a positive definite matrix, according to Sherman-Morrison equation, the first part of this new upper bound can be rewritten as:

$$\alpha_{x_a}^{t+1} \|\hat{a}_{y_a}^t \circ \hat{s}_t^t\|_{(A_{x_a}^{t+1})^{-1}}^2 = \frac{\alpha_{x_a}^{t+1} \|\hat{a}_{y_a}^t \circ \hat{s}_t^t\|_{(A_{x_a}^t)^{-1}}^2}{1 + \|\hat{a}_{y_a}^t \circ \hat{s}_t^t\|_{(A_{x_a}^t)^{-1}}^2},$$

where we use $A_{x_a}^t$ to replace $A_{x_a}^{t+1}$. The same technique can be applied to the other parts, and thus, it is easy to derive the upper bound of the estimation error:

$$E_A(t + 1) \leq \frac{\alpha_1 M}{\sqrt{1 + M^2}} + \frac{\alpha_2 N}{\sqrt{1 + N^2}} + \alpha_1 G + \alpha_2 H + C$$

$$E_O(t + 1) \leq \frac{\alpha_1 G}{\sqrt{1 + G^2}} + \frac{\alpha_2 H}{\sqrt{1 + H^2}} + \alpha_1 M + \alpha_2 N + C$$

$$\text{where } M = \|\hat{a}_{y_a}^t \circ \hat{s}_t^t\|_{(A_{x_a}^t)^{-1}}, N = \|\hat{h}_{x_a}^t \circ \hat{s}_t^t\|_{(C_{y_a}^t)^{-1}}$$

$$G = \|\hat{a}_{x_o}^t \circ \hat{s}_t^t\|_{(A_{x_o}^t)^{-1}}, H = \|\hat{h}_{y_o}^t \circ \hat{s}_t^t\|_{(C_{y_o}^t)^{-1}}$$

Therefore, according to the selection strategy defined in ActSense:

$$\alpha_1 M + \alpha_2 N \geq \alpha_1 G + \alpha_2 H, \text{ with } M \geq G \text{ and } N \geq H,$$

we can prove that with high probability, the upper bound of estimation error of ActSense is smaller than the one generated by any other selections. $\square$

As $E_O(t + 1)$ represents the predicted error generated by any other selections, this conclusion means our method can achieve the best prediction error reduction among all possible selections, or at least no worse than any one of them.

4 EMPIRICAL EVALUATION

We use the Dataport dataset [29] for evaluation. It is the largest public residential home energy dataset, which contains the appliance-level and household aggregate energy consumption sampled every minute from 2012 to 2018. While Dataport contains energy data from various cities in the U.S., we only focus on the data collected from Austin as it contains the largest set of homes (i.e., 534 homes) from a single region. We filter out the appliances with poor data quality (large proportion of missing values). We get 4 different datasets for each year between 2014 to 2017 containing 53, 93, 73, and 44 homes respectively and six appliances: air-conditioning (HVAC), fridge, washing machine, furnace, microwave and dishwasher. On this selected data set, we reconstruct the aggregate reading by the sum of the selected appliances [16, 24].

An important reason for choosing these appliances is that they represent a wide variety of household energy consumption patterns. For example, season dependent (HVAC) v.s., season independent




Figure 2: Energy breakdown for Austin in 2015 and 2016.

(dishwasher), background (fridge) v.s., interactive (microwave), etc. Figure 2 shows the appliance level energy consumption pattern in each month in Austin across two adjacent years. The usage patterns across months are quite similar. For example, higher consumption during summer (dominated by HVAC) and lower in winter. The energy consumption for the remaining two years while not shown is fairly similar to these results as well. This observation supports our previous assumption that the season pattern is similar across years within one geo-region, and thus, we could re-use the season patterns learned from historical information to guide future projection.

Reproducibility Our entire codebase, baselines, analysis and experiments can be found on Github ².

4.1 Baseline Algorithms

In our evaluation, we use the following baselines.

- **Random:** performs CP decomposition with ALS and selects L <home, appliance> pairs uniformly random from the candidate pool. Random sampling has often been used as a baseline in prior active learning literature and also represents a feasible mechanism via which utility companies actually deploy sensors in practice.
- **QBC:** performs CP decomposition with ALS and select pairs with Query-by-Committee (QBC) strategy. The QBC framework quantifies the prediction uncertainty based on the level of disagreement among a set of trained models. Similar to previous work [10], we perform tensor factorization using ALS with different settings of the rank parameter to form the committee. The estimation uncertainty on each <home, appliance> pair is computed by the variance across the estimates of the committee members.
- **VBV:** performs CP decomposition with variational Bayesian inference and selects pairs according to the variance of each estimation. Variational Bayesian (VB) [9] is one of the most popularly used methods for uncertainty assessment. Prior work [37] proposed VB for active completion in a matrix setting. In this paper, we extend their setting to tensor completion according to the work by Zhao et. al [41]. At the end of each month, VB estimates the posterior distributions of latent factors and selects the <home, appliance> pairs with the highest posterior variance in the energy estimation.

We should note that while we discussed a few other active matrix completion methods in Section 2, we did not use them as baselines due to their poor scalability in the energy breakdown scenario. For example, the graphical lasso algorithm [10] for matrix completion needed at least 50% of the total observations to estimate the model,

and assumed the dependency between the missing values and observations follows a Gaussian distribution. A prior VB based matrix completion work [36] proposed to use the mutual information as the selection criteria. However, they model the mutual information between rows and between columns in the matrix, and select entire row or column separately. This would require us to instrument one appliance in all homes, which is clearly infeasible in practice.

4.2 Evaluation Metric & Setup

Based on prior literature [6, 8], we evaluate the performance of energy breakdown with root mean square error (RMSE) between the predicted appliance energy and the ground-truth on the test set. In month t , the ground-truth and estimated energy for appliance j at home i are denoted as e_{ijt} and $\hat{e}_{ijt}$. The RMSE of appliance j at month t is given as,

$$RMSE(j, t) = \sqrt{\frac{1}{M} \sum_i (e_{ijt} - \hat{e}_{ijt})^2}$$

where M represents the total number of homes in the test set. For month t , we use the Mean RMSE across appliances to measure a model's prediction accuracy:

$$Mean\ RMSE(t) = \frac{\sum_j RMSE(j, t)}{N}$$


where N represents the number of appliances. Lower Mean RMSE indicates better energy breakdown performance.

We use 5-fold cross-validation across homes in all our experiments. The last 20% of the train set in each fold is reserved for validation purpose. For each baseline and our method, the optimal parameters (e.g., rank of the tensor) are learned via an exhaustive grid search on the validation set. We first fix α_i^t and α_j^t to 0.1. For tensor decomposition, we vary the rank of the latent dimensions from 1 to 4. With limited observations, the model could overfit, especially in the first several rounds of optimization. Thus, a larger regularization coefficient should be applied to avoid it. We choose the λ values from {5000, 8000, 10000}. For the kernel function used in our uncertainty integration, we vary the window size σ from {1, 3, 6, 12}.


4.3 Experiment Results

4.3.1 Quality of Energy Breakdown. In this experiment, we compare the performance of ActSense and the baselines in monthly error, which represents the instant energy breakdown performance. We fixed the number of selections for each month as $L = 5$, so that for at the end of each year of our evaluation, we will have 10.75% to 22.73% <home,appliance> pairs instrumented across the selected four years. Figure 3 shows the quality of energy breakdown in Austin, 2015. Mean RMSE across appliances for each month is reported in Figure 3 (a). We can observe that our uncertainty based selection performs favourably compared to the baselines. We also plot the relative improvement compared to the random baseline in each month in Figure 3 (b). We can observe that our approach achieves the highest improvement compared to other baselines, especially in June, where the improvement is up to 35.06%. To put this improvement into context, June and other summer months typically have the highest energy usage, where the scope and potential benefits of energy breakdown are also high.


²<https://github.com/yilingjia/ActSense>



(a) Mean RMSE performance across months.



(b) Relative improvement compared with random method.

Figure 3: Our approach ActSense gives the best energy breakdown performance across all months for the year 2015.**Figure 4: Selection ratio of appliances, Austin, 2015.**

We now discuss why some baselines failed to provide an accurate energy breakdown. Random selection uniformly selects the <home, appliance> pairs, ignoring the difference in their informativeness. For example, microwave is a season independent appliance, whose usage pattern should be simpler than season dependent appliances, such as HVAC. As random selection treats all the appliances equally, additional selections made on those well-learned appliances are wasted, while the latent factors for appliances with complicated usage patterns are not well modeled. Figure 4 shows the selection ratio of appliances across different methods. All of them except Random selected more on HVAC or furnace, which are more difficult to model as they are highly season dependent. While VBV imposes the same tensor structure, it uses a different parameter estimation procedure compared to the other methods. VBV gives worse energy breakdown quality than other approaches due to its poor performance in parameter estimation. VBV estimates the

Table 1: Relative improvement comparing to Random.

	Maximum Improvement (%)			Average Improvement (%)		
	QBC	ActSense	VBV	QBC	ActSense	VBV
2014	12.30	29.71	-2.21	1.98	7.07	-76.85
2015	23.09	35.06	-2.07	5.58	11.88	-41.71
2016	22.94	29.84	-14.01	9.38	14.66	-62.97
2017	7.42	28.76	0.32	4.48	12.07	-40.82

distribution of each latent factor, including both mean and variance. Its model complexity is higher, and thus it can easily overfit. From its selections, we can find that it concentrated on HVAC and seldom selected other appliances. Though HVAC needs more observations to get a good estimation, VBV's selection overfits to HVAC and give worse energy performance for other appliances. Among the baselines, QBC gives a reasonable energy breakdown quality. However, as discussed before, QBC will select the pairs which have the highest variance among the committee members. This selection strategy would perform well when data has a similar scale. But in energy breakdown, the energy readings across appliances and homes are heavily imbalanced (shown in Figure 2). The selection is easily dominated by the appliances with high energy consumption, such as HVAC and furnace as they definitely have a higher variance than the "minor" appliances which in general consume lower energy, in models' predictions. Our uncertainty based active selection strategy overcomes the problems discussed before. First, it balances the selections among the appliances. Though most of the selections are for HVAC and furnace, the other "minor" appliances are not ignored. Every appliance has a chance to be selected and modeled. Second, as we incorporate the temporal information into our selection strategy, ActSense can foresee the change in future and make up for the mistakes in the past. Intuitively, we should have more observations from HVAC in summer to accommodate its changing usage pattern. And the home and appliance pairs that the current model is still uncertain about based on historical observation will be selected with high probability. Figure 4 shows that our method selects most observations of HVAC in May, while QBC selects most HVAC starting from June to September. This indicates our model foresees the future change and prepares for the sensor installation in advance. Therefore, its performance in summer is much better than the other baselines. Moreover, our approach also integrates the uncertainty from the previous months to make up for the mistakes. We can see that after May, ActSense tends to select more on the other appliances so that we can have a good model for HVAC without sacrificing the performance of others.

Due to the space limit, we only report the results in other years with the best and average relative improvement across months in Table 1; and the detailed results are provided in our supplementary material. We can observe ActSense gives encouraging improvement over the baselines in general across these four years' evaluation.

4.3.2 Budget size. As discussed before, sensor installation is expensive, and the goal of active sensor deployment is to maximize the energy breakdown accuracy while minimizing the cost of instrumentation. In this experiment, we want to explore how the budget size affects the energy breakdown performance in each method. We fix the hyper-parameters with the best set found in the grid search, and only vary the number of <home, appliance> pairs selected in




Figure 5: Performance v.s., the number of selections in each month, Austin, year 2015.

Table 2: Relative Improvement comparing to Random with different uncertainty estimation.

Uncertainty Estimation	Maximum	Mean
Current	34.38%	11.48%
Current + Future	34.89%	11.82%
History + Current + Future	35.06%	11.88%

each month from 1 to 20. With such a setting, the total number of sensors installed during one year will vary from 11 to 220. With total 93 homes and 6 appliances in the dataset from Austin in the year of 2015, this setting makes the installation ratio vary from 1.97% to 39.43%. We evaluate each method by its average energy breakdown performance in one year by,

$$\text{Year RMSE} = \frac{\sum_{t=1}^{12} \text{Mean RMSE}(t)}{12}$$

Figure 5 shows our approach gives consistently better performance and faster convergence with an increasing size of the sensor installation budget. For a more practical comparison, with a fixed target of energy breakdown, say setting Year RMSE to 50, shown in the figure, our approach only needs 3 new observations per month, while QBC needs 8 and Random needs 10 new observations. This clearly demonstrates our solution’s advantage in minimizing the cost of sensor deployment for improving energy breakdown quality.

4.3.3 Temporal information. In this section, we analyze the contribution of the temporal information incorporated into our selection strategy. In the experiment, we compare our proposed method with different uncertainty estimations: 1) Current: selects pairs based on the uncertainty on current month only; 2) Current + Future: combines the current uncertainty and the future uncertainty; 3) History + Current + Future: integrates these three types of uncertainty estimation as in our Algorithm 1.

The results in Table 2 show the best and average relative improvement of different uncertainty estimation techniques over Random across months. It can be seen that both future and historical estimated uncertainty improve the energy breakdown performance relative to Random. Furthermore, we can observe that the major contribution comes from future projection. Figure 6 shows the selection ratio of appliances with different uncertainty estimation methods. We can notice, 1) with future projection, our approach




Figure 6: Selection ratio of appliances for different uncertainty measurement.

can prepare for future changes in advance, for example, it tends to query more observations of HVAC in April and May; 2) with historical uncertainty, the model better balances between the major and minor appliances. With historical information, the algorithm can detect mistakes made earlier, and query more to make up for the same. For example, compared to the other two uncertainty measurements, the model with historical uncertainty tends to select more washing machine and HVAC at the end of year, instead of only focusing on furnace. The results in Table 2 indicates this knowledge retrospect helps for better energy breakdown performance.

5 DISCUSSIONS

In addition to the promising empirical and theoretical results obtained by our proposed solution, there are a few limitations of our current work which we plan to address in the future.

• **Heterogeneous cost:** In our current setting, we assume that the cost of instrumenting different appliances are the same. However, in practice, the costs may vary due to different difficulty of instrumentation and the labor cost. The key of solving the problem is to balance the uncertainty reduction and the cost of sensor installation during the active deployment. Therefore, the selection strategy should also take the heterogeneous costs into consideration, e.g., solving a budget constrained optimization problem.

• **Dynamic installation under budget constraints:** Our current setting assumes the number of sensors installed in each month is fixed, denoted with L in the paper. However, another practical setting for sensor deployment is to consider a fixed total number of sensors. This introduces another question - how to distribute sensor deployment across months? As discussed earlier, the appliance energy consumption is highly related with season. This indicates that the complexity of usage patterns across months is different. Thus, in order to get a good estimation of energy breakdown, the number of observations required for each month should be different. Similar to the way we leveraged temporal information in our current method, we can estimate the distribution of uncertainties in future months, and dynamically change the number of sensor installation in each month according to the total budget constraint.

• **Transferable active learning:** Our current approach only works for a single region. We assume that within one region, the structure across homes is similar due to the common design and season pattern. However, the available energy data across geo-regions is imbalanced. It is hard to generate a good energy breakdown model for a region with limited available data. We can model homogeneous or region-independent factors about the appliances, and the

similarity of homes from source region and target region to actively install sensors for the target region.

6 CONCLUSION

Active collaborative sensing for sensor deployment aims to balance the performance of energy breakdown and the cost of instrumentation. The main challenge is to select the most representative or informative observations for the non-instrumented homes. To the best of our knowledge, this is the first work addressing the active sensing problem in energy breakdown. In our work, we quantify the uncertainty in the parameter estimation process for query selection. We also integrate temporal information to retrospect history and foresee future. Our theoretical analysis and empirical evaluation prove that our approach performs favourably compared to the state-of-the-art. We believe that our method has the potential to create a paradigm shift in sensor deployments and generate positive contribution to global energy saving.

7 ACKNOWLEDGEMENTS

We want to thank the reviewers for their insightful comments. This work is based upon work supported by National Science Foundation under grant CNS-1646501, IIS-1553568, IIS-1718216.

REFERENCES

- [1] K Carrie Armel, Abhay Gupta, Gireesh Shrimali, and Adrian Albert. 2013. Is disaggregation the holy grail of energy efficiency? The case of electricity. *Energy Policy* 52 (2013), 213–234. <https://doi.org/10.1016/j.enpol.2012.08.062>
- [2] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47, 2-3 (2002), 235–256.
- [3] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. 1995. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *focs*. IEEE, 322.
- [4] Joseph E Banta, LR Wong, Christophe Dumont, and Mongi A Abidi. 2000. A next-best-view system for autonomous 3-D object reconstruction. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 30, 5 (2000), 589–598.
- [5] Nipun Batra, Rishi Bajjal, Amarjeet Singh, and Kamin Whitehouse. 2015. How good is good enough? Re-evaluating the bar for energy disaggregation. *arXiv preprint arXiv:1510.08713* (2015).
- [6] Nipun Batra, Yiling Jia, Hongning Wang, and Kamin Whitehouse. 2018. Transferring Decomposed Tensors for Scalable Energy Breakdown across Regions. *AAAI* (2018).
- [7] Nipun Batra, Amarjeet Singh, and Kamin Whitehouse. 2017. *Systems and analytical techniques towards practical energy breakdown for homes*. Ph.D. Dissertation. IIT-Delhi.
- [8] Nipun Batra, Hongning Wang, Amarjeet Singh, and Kamin Whitehouse. 2017. Matrix Factorisation for Scalable Energy Breakdown.. In *AAAI*. 4467–4473.
- [9] Matthew James Beal et al. 2003. *Variational algorithms for approximate Bayesian inference*. university of London London.
- [10] Shayok Chakraborty, Jiayu Zhou, Vineeth Balasubramanian, Sethuraman Panchanathan, Ian Davidson, and Jieping Ye. 2013. Active matrix completion. In *Data Mining (ICDM), 2013 IEEE 13th International Conference on*. IEEE, 81–90.
- [11] SY Chen and YF Li. 2005. Vision sensor planning for 3-D model acquisition. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 35, 5 (2005), 894–904.
- [12] Shengyong Chen, Youfu Li, and Ngai Ming Kwok. 2011. Active vision in robotic systems: A survey of recent developments. *International Journal of Robotics Research* 30, 11 (2011), 1343–1377.
- [13] David A Cohn, Zoubin Ghahramani, and Michael I Jordan. 1996. Active learning with statistical models. *Journal of artificial intelligence research* 4 (1996), 129–145.
- [14] Samuel DeBruin, Branden Ghena, Ye-Sheng Kuo, and Prabal Dutta. 2015. Powerblade: A low-profile, true-power, plug-through energy meter. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*. ACM, 17–29.
- [15] Anthony Faustine, Nerey Henry Mvungi, Shubi Kaijage, and Kisangiri Michael. 2017. A survey on non-intrusive load monitoring methodologies and techniques for energy disaggregation problem. *arXiv preprint arXiv:1703.00785* (2017).
- [16] Zoubin Ghahramani and Michael I Jordan. 1997. Factorial hidden Markov models. *Machine learning* 29, 2-3 (1997).
- [17] George William Hart. 1992. Nonintrusive appliance load monitoring. *Proc. IEEE* 80, 12 (1992), 1870–1891. <https://doi.org/10.1109/5.192069>
- [18] Neil Houlsby, José Miguel Hernández-Lobato, and Zoubin Ghahramani. 2014. Cold-start active learning with robust ordinal matrix factorization. In *ICML*. 766–774.
- [19] Xiaofan Jiang, Stephen Dawson-Haggerty, Prabal Dutta, and David Culler. 2009. Design and implementation of a high-fidelity ac metering network. In *2009 International Conference on Information Processing in Sensor Networks*. IEEE, 253–264.
- [20] Srinivas Katipamula and Michael R Brambley. 2005. Review article: methods for fault detection, diagnostics, and prognostics for building systems—A review. *HVAC Research* (2005), 3–25.
- [21] Jaya Kawale, Hung H Bui, Branislav Kveton, Long Tran-Thanh, and Sanjay Chawla. 2015. Efficient Thompson Sampling for Online Matrix-Factorization Recommendation. In *NIPS*. 1297–1305.
- [22] Jack Kelly, Nipun Batra, Oliver Parson, Haimonti Dutta, William Knottenbelt, Alex Rogers, Amarjeet Singh, and Mani Srivastava. 2014. Nilmk v0. 2: a non-intrusive load monitoring toolkit for large scale data sets: demo abstract. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*. ACM, 182–183.
- [23] Jack Kelly and William Knottenbelt. 2016. Does disaggregated electricity feedback reduce domestic electricity consumption? A systematic review of the literature. *arXiv preprint arXiv:1605.00962* (2016).
- [24] J. Z. Kolter, S. Batra, and A. Y. Ng. 2010. Energy Disaggregation via Discriminative Sparse Coding. In *NIPS 2010*. Vancouver, BC, Canada.
- [25] J. Z. Kolter and T. Jaakkola. 2012. Approximate Inference in Additive Factorial HMMs with Application to Energy Disaggregation. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*. La Palma, Canary Islands.
- [26] Kiriakos N Kutulakos and Charles R Dyer. 1995. Global surface reconstruction by purposive control of observer motion. *Artificial Intelligence* 78, 1 (1995), 147–177.
- [27] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th WWW*. ACM, 661–670.
- [28] YF Li and ZG Liu. 2005. Information entropy-based viewpoint planning for 3-D object reconstruction. *IEEE Transactions on Robotics* 21, 3 (2005), 324–337.
- [29] Oliver Parson, Grant Fisher, April Hersey, Nipun Batra, Jack Kelly, Amarjeet Singh, William Knottenbelt, and Alex Rogers. 2015. Dataport and NILMTK: A building data set designed for non-intrusive load monitoring. In *GlobalSIP 2015*. IEEE.
- [30] Oliver Parson, Siddhartha Ghosh, Mark J Weal, and Alex Rogers. 2012. Non-Intrusive Load Monitoring Using Prior Models of General Appliance Types.. In *AAAI*.
- [31] Luis Pérez-Lombard, José Ortiz, and Christine Pout. 2008. A review on buildings energy consumption information. *Energy and buildings* 40, 3 (2008), 394–398.
- [32] Richard Pito. 1999. A solution to the next best view problem for automated surface acquisition. *IEEE PAMI* 21, 10 (1999), 1016–1030.
- [33] Burr Settles. 2012. Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 6, 1 (2012), 1–114.
- [34] Huijuan Shao, Manish Marwah, and Naren Ramakrishnan. 2013. A Temporal Motif Mining Approach to Unsupervised Energy Disaggregation: Applications to Residential and Commercial Buildings.. In *AAAI*.
- [35] Ksenia Shubina and John K Tsotsos. 2010. Visual search for an object in a 3D environment using a mobile robot. *Computer Vision and Image Understanding* 114, 5 (2010), 535–547.
- [36] Jorge Silva and Lawrence Carin. 2012. Active learning for online bayesian matrix factorization. In *Proceedings of the 18th ACM SIGKDD*. ACM, 325–333.
- [37] Dougal J Sutherland, Barnabás Póczos, and Jeff Schneider. 2013. Active learning and search on low-rank matrices. In *Proceedings of the 19th ACM SIGKDD*. ACM, 212–220.
- [38] Simon Tong and Daphne Koller. 2001. Support vector machine active learning with applications to text classification. *JMLR* 2, Nov (2001), 45–66.
- [39] André Uschmajew. 2012. Local convergence of the alternating least squares algorithm for canonical tensor approximation. *SIAM J. Matrix Anal. Appl.* 33, 2 (2012), 639–652.
- [40] Huazheng Wang, Qingyun Wu, and Hongning Wang. 2017. Factorization Bandits for Interactive Recommendation.. In *AAAI*. 2695–2702.
- [41] Qibin Zhao, Guoxu Zhou, Liqing Zhang, Andrzej Cichocki, and Shun-Ichi Amari. 2016. Bayesian robust tensor factorization for incomplete multiway data. *IEEE transactions on neural networks and learning systems* 27, 4 (2016), 736–748.
- [42] Ahmed Zoha, Alexander Gluhak, Muhammad Ali Imran, and Sutharshan Rajasegarar. 2012. Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey. *Sensors* 12, 12 (2012), 16838–16866.