Inferring the Behavior of Distributed Energy Resources with Online Learning

Gregory S. Ledva, Laura Balzano, Johanna L. Mathieu

Abstract— **In this paper, we apply an emerging method, online learning with dynamics, to deduce properties of distributed energy resources (DERs) from coarse measurements, e.g., measurements taken at distribution substations, rather than household-level measurements. Reduced sensing requirements can lower infrastructure costs associated with reliably incorporating DERs into the distribution network. We specifically investigate whether dynamic mirror descent (DMD), an online learning algorithm, can determine the real-time controllable demand served by a distribution feeder using feeder-level active power demand measurements.**

In our scenario, DMD incorporates various controllable demand and uncontrollable demand models to generate real-time controllable demand estimates. In a realistic scenario, these estimates have an RMS error of 8.34% of the average controllable demand, which improves to 5.53% by incorporating more accurate models. We propose topics for additional work in modeling, system identification, and the DMD algorithm itself that could improve the RMS errors.

I. INTRODUCTION

Distributed energy resources (DERs), e.g., demand responsive loads and residential solar generation, are becoming more prevalent within distribution networks [1], [2]. Sensing capabilities enabled by technologies such as residential smart meters [3] are also becoming more prevalent in power systems, but electric utilities still lack an accurate real-time picture of the behavior of DERs. Such information could help improve power system reliability, economic efficiency, and environmental impact. Realistically, we do not need perfect information about every device connected to the system in order to achieve these benefits. This work explores using online methods to infer the behavior of DERs from coarse measurements taken from distribution substations rather than more detailed household-level measurements. Effective inference methods would reduce the need for sensing and communications infrastructure, reducing the costs of incorporating DERs.

Extensive literature exists for non-intrusive load monitoring (NILM), for example [4], or energy disaggregation, for example [3], which we will collectively refer to as NILM. NILM seeks to infer individual appliance or load behavior, usually for less than 20 loads, from a single aggregate power measurement signal sampled at high frequency (e.g., 10-100 kHz) at a household main. These algorithms usually are not solved online because the goal is long-term energy efficiency decisions such as identification and replacement of faulty appliances and/or load research. Many of the existing techniques for NILM rely on assumptions, for example, that only one device can change state at a time [5] or that step changes can been seen in the aggregate data [6], which enables the use of change detection algorithms.

In contrast, we want to infer information in real-time, i.e., as measurements arrive. The measurements could be a feeder's demand served, voltage, etc., which are impacted by large numbers of devices (thousands to tens of thousands). As a result, typical NILM assumptions are invalid and the corresponding NILM approaches are inapplicable. Another difference from NILM is that rather than determining device-level behavior, we seek to determine the behavior of classes of devices. For example, in real-time, a demand response aggregator cares about the total controllable demand served by a feeder. However, after service delivery, the aggregator does care about device-level information for auditing or compensation.

In this paper, we investigate the application of an online learning algorithm, dynamic mirror descent (DMD) [7], to determine DERs behavior in real-time. Specifically, we use DMD to deduce the controllable demand component of feeder-level active power measurements as they arrive. We detail several models used to capture the behavior of the controllable and uncontrollable demand, and we propose several areas where these models could be improved for our application. Finally, we present several case studies showing the capabilities and limitations of DMD within the proposed scenario.

This paper is organized as follows: Section II presents the problem setting. Section III constructs the physical representation of the system. Section IV describes the online learning algorithm, and Section V describes the models used within this algorithm. Section VI presents several case studies and their results, and Section VII provides concluding remarks.

II. PROBLEM SETTING

In this paper, we consider a situation where an entity, e.g., a demand response aggregator, has access to real-time active power measurements corresponding to the load, or demand, served by a distribution feeder. The power measurements consist of time-averaged active power draw over one minute intervals. A portion of each active power demand measurement corresponds to the power usage of controllable appliances, i.e., demand that the aggregator can manipulate. The remaining portion of each measurement corresponds to the uncontrollable demand, i.e., the demand that cannot be manipulated. The aggregator wants to deduce the controllable demand portion of the total demand measurements as the

The authors are with the Department of Electrical Engineering at the University of Michigan. This research was funded by NSF Grant #ECCS-1508943.

measurements arrive. Knowledge of the real-time controllable demand has several uses: i) to bid demand response capacity into ancillary services markets [8], [9], or ii) as the feedback control signal within a load control algorithm [10]–[14]. While this knowledge enables feedback control, in this work we do not incorporate any control action that manipulates the controllable loads.

We apply an online learning algorithm, dynamic mirror descent (DMD) [7], to achieve the goal of estimating the controllable demand in real-time. DMD uses a bank of models, some for the controllable demand and some for the uncontrollable demand, to predict the former. The algorithm forms overall demand predictions from various combinations of the controllable and uncontrollable demand models. These model combinations are weighted based on their recent predictions' accuracy – better prediction-measurement matching leads to larger weighting and more influence in the overall estimate. DMD's controllable demand estimate is then the controllable demand component of this weighted combination of the various model predictions.

Within this work, we assume that the controllable demand corresponds to the active power usage of a population of thousands of air conditioners. We assume the air conditioners are all connected to a single distribution feeder. Each air conditioner cycles between drawing power (the on mode) and not drawing power (the off mode) to maintain a household's indoor temperature within a dead-band around a user-defined temperature setting. We also assume that smart meters measure the indoor temperature and on/off mode of each air conditioner on the time-scale of seconds. We assume this information is not available in real-time, but historical data, e.g., from a previous day, are available.

In the following section, we describe the general simulation setting and the construction of the plant, which is the representation of the physical system within our simulations.

III. PLANT CONSTRUCTION

In this section, we detail the construction of a 24-hour signal representing one day of the distribution feeder's measured power values, which are the summation of a controllable demand signal and an uncontrollable demand signal. We generate the uncontrollable demand signal using a feeder from GridLAB-D's feeder taxonomy [15] and household power usage data from the Pecan Street Inc. Dataport [16]. We generate the controllable demand signal from a set of models that simulate the power draw of individual air conditioners.

Figure 1 shows the time series corresponding to the controllable demand, denoted y_t^{c}, and uncontrollable demand, denoted y_t^{uc}. The measured demand time series y_t is the sum of these two signals, $y_t = y_t^{\text{c}} + y_t^{\text{uc}}$, and the 1 minute sampling intervals for the Dataport's household data dictates the time-step of our simulation. The following sections detail the construction of the uncontrollable demand signal and the controllable demand signal.

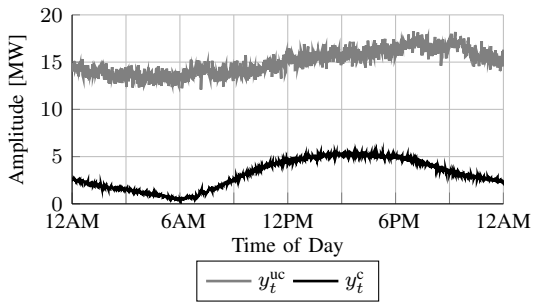

Fig. 1. Time series of the controllable and uncontrollable demand components of the measured signal.

A. Uncontrollable Demand Signal

The active power draw of the loads on GridLAB-D feeder R3-12.47-1 determines the uncontrollable demand signal's average value over the day. The chosen feeder contains only residential and commercial loads, and we assume the commercial load served by the feeder is constant over the day. To add some variability to the commercial load, we add zero-mean, normally distributed demand with a standard deviation of 5% of the mean commercial load.

The residential loads served by the feeder provide the average value of the uncontrollable household load. We construct the uncontrollable household load signal from the Dataport's historical household active power demand data. The data used to construct the uncontrollable residential load was that of single family homes in Austin, Texas on Monday, August 3, 2015, and we removed any air conditioning loads from the data. The individual houses' usage signals for the day were randomly drawn with replacement and added to the uncontrollable residential signal until the uncontrollable residential signal's mean matched that dictated by the feeder model. Finally, the commercial and residential components were added to generate the overall uncontrollable demand signal. Note that we do not model power lines or power flows within this work.

B. Controllable Demand Signal

A population, or set, of hybrid models, i.e., models containing continuous and discrete states, generates the controllable demand signal. Each hybrid model captures the heat transfer driving the on/off cycling of an individual air conditioner. The total controllable demand at time t, y_t^{c}, is the sum of the individual hybrid models' power usage.

The number of models within the population, N^{ac}, equals the number of houses with air conditioners used to generate the uncontrollable demand signal, and we model an individual air conditioner's power usage with the equivalent thermal parameter (ETP) model [17]. The set of ETP models are $\mathcal{N}^{\text{ac}} = \{1, 2, \ldots, N^{\text{ac}}\}$, and we index them with $i \in \mathcal{N}^{\text{ac}}$.

The ETP model contains a vector of continuous states $\theta_t^i = \begin{bmatrix}\theta_t^{a,i} & \theta_t^{m,i}\end{bmatrix}^T$ that capture the house's internal temperatures and a discrete state m_t^i that captures whether the appliance is drawing power. The $\theta_t^{m,i}$ value is the house's internal mass temperature, and $\theta_t^{a,i}$ is the house's internal air temperature. Table I summarizes the parameters used within the model and

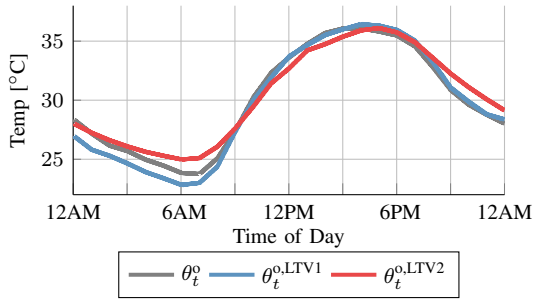

Fig. 2. Time series plots comparing θ_t^{o}, the outdoor temperature trajectory used to generate the controllable demand signal, with $\theta_t^{\mathrm{o,LTV1}}$ and $\theta_t^{\mathrm{o,LTV2}}$, two alternative trajectories used to generate the models in Section V-B.3.

their values, where $[\alpha, \beta]$ indicates a uniform distribution. We generate the N^{ac} ETP models by randomly sampling the parameter distributions. The nominal parameters for the parameter distributions are based on [18].

The vector $d_t^i = \begin{bmatrix}\theta_t^{\mathrm{o}} & Q^{\mathrm{a},i} & Q^{\mathrm{m},i}\end{bmatrix}^T$ collects the environmental heat sources that influence the on/off cycling of the air conditioners. The values $Q^{\mathrm{a},i}$ and $Q^{\mathrm{m},i}$ capture heating from appliances and occupants within the house as well as solar irradiance. In this work, they are assumed constant. The outdoor temperature is θ_t^{o}, and its trajectory for August 3 is shown in Fig. 2. The Dataport weather data is sampled hourly, and we interpolated to the necessary time-step. Future work should investigate time-varying values for $Q^{\mathrm{a},i}$ and $Q^{\mathrm{m},i}$ that depend on weather conditions, time of day, and uncontrollable load. These values are a function of the uncontrollable load because the uncontrollable load includes uncontrollable appliances, e.g., computers, that induce heat gains.

To create the dynamical models, we first form continuous time matrices from the parameters sampled from Table I

$$A^{\mathrm{c},i} = \begin{bmatrix} -\left(U^{\mathrm{a},i} + U^{\mathrm{m},i}\right)/\Lambda^{\mathrm{a},i} & U^{\mathrm{m},i}/\Lambda^{\mathrm{a},i} \\ U^{\mathrm{m},i}/\Lambda^{\mathrm{m},i} & -U^{\mathrm{m},i}/\Lambda^{\mathrm{m},i} \end{bmatrix} \quad (1)$$

$$B^{\mathrm{c},i} = \begin{bmatrix} Q^{\mathrm{h},i}/\Lambda^{\mathrm{a},i} & 0 \end{bmatrix}^T \quad (2)$$

$$E^{\mathrm{c},i} = \begin{bmatrix} U^{\mathrm{a},i}/\Lambda^{\mathrm{a},i} & 1/\Lambda^{\mathrm{a},i} & 0 \\ 0 & 0 & 1/\Lambda^{\mathrm{m},i} \end{bmatrix}. \quad (3)$$

We form the discrete-time matrices using [19, p. 315] and use them to update θ_t^i and m_t^i

$$\theta_{t+1}^i = A^i\theta_t^i + B^i m_t^i + E^i d_t^i \qquad i \in \mathcal{N}^{\mathrm{ac}} \quad (4)$$

$$m_{t+1}^i = \begin{cases} 0 & \text{if } \theta_{t+1}^{a,i} < \theta^{\mathrm{set},i} - \theta^{\mathrm{db},i}/2 \\ 1 & \text{if } \theta_{t+1}^{a,i} > \theta^{\mathrm{set},i} + \theta^{\mathrm{db},i}/2 \\ m_t^i & \text{otherwise} \end{cases} \quad i \in \mathcal{N}^{\mathrm{ac}}. \quad (5)$$

The first equation updates the internal temperatures. The second equation updates the on/off mode based on whether or not the air temperature has exited the allowable temperature range. The resulting power draw is $P_t^i = (|Q^{\mathrm{h},i}|\ m_t^i)/\eta^i$ where heat is removed, i.e., $Q^{\mathrm{h},i} < 0$, for cooling appliances.

The controllable demand value at each time-step is $y_t^{\mathrm{c}} = \sum_{i=1}^{N^{\mathrm{ac}}} P_t^i$. Note that the ETP models operate on time-steps of two seconds while the demand signals have one minute measurement intervals. The y_t^{c} within each one minute interval are averaged to form the controllable demand component of the measurement.

TABLE I

AIR CONDITIONER MODEL PARAMETERS [18]

Parameter	Description	Value
Δt	Time-Step Duration [s]	2
θ^{set}	Temperature Set-Point [°C]	[24, 26]
θ^{db}	Temperature Dead-band [°C]	[2.0, 2.5]
U^{m}	Envelope Conductance [$\frac{\mathrm{kW}}{^\circ\mathrm{C}}$]	[0.84, 1.14]
U^{a}	Internal Conductance [$\frac{\mathrm{kW}}{^\circ\mathrm{C}}$]	[0.2, 0.27]
Λ^{m}	Mass Heat Capacitance [$\frac{\mathrm{kWh}}{^\circ\mathrm{C}}$]	[4.48, 6.07]
Λ^{a}	Air Heat Capacitance [$\frac{\mathrm{kWh}}{^\circ\mathrm{C}}$]	[0.16, 0.21]
Q^{m}	Internal Mass Heat Gain [kW]	0.5
Q^{a}	Internal Air Heat Gain [kW]	0.5
Q^{h}	Appliance Heat Transfer [kW]	[-17.7, -13.1]
η	Coefficient of Performance [-]	3

IV. DYNAMIC MIRROR DESCENT

In this section, we detail DMD [7] and adapt it to the scenario under consideration. In our scenario, DMD seeks to estimate the controllable demand component, y_t^{c}, of the total demand measurements, $y_t = y_t^{\mathrm{c}} + y_t^{\mathrm{uc}}$, as they arrive. To do this, we define the set of N^{mdl} models used within DMD as $\mathcal{N}^{\mathrm{mdl}} = \{1, \ldots, N^{\mathrm{mdl}}\}$.

DMD estimates y_t^{c} using two general processes. The first process forms predictions $\widehat{y}_t^{\mathrm{c},i}$ and $\widehat{y}_t^{\mathrm{uc},i}$ for each model $i \in \mathcal{N}^{\mathrm{mdl}}$, and then uses the corresponding model to advance the predictions in time. The second process determines a weight associated with each model and forms an overall estimate, $\widehat{y}_t = \widehat{y}_t^{\mathrm{c}} + \widehat{y}_t^{\mathrm{uc}}$, using a weighted combination of the $\widehat{y}_t^{\mathrm{c},i}$ and $\widehat{y}_t^{\mathrm{uc},i}$ values. The value of interest to the aggregator is $\widehat{y}_t^{\mathrm{c}}$, the controllable demand component of this weighted combination. We first detail the prediction process, then the weighting process that forms the estimate.

We associate with each model a prediction $\widehat{\theta}_t^i = \begin{bmatrix}\widehat{y}_t^{\mathrm{c},i} & \widehat{y}_t^{\mathrm{uc},i}\end{bmatrix}^T$, a modified prediction $\widetilde{\theta}_t^i = \begin{bmatrix}\widetilde{y}_t^{\mathrm{c},i} & \widetilde{y}_t^{\mathrm{uc},i}\end{bmatrix}^T$, and a specified model of arbitrary structure $\Phi_t^i(\cdot)$. Note that in our scenario, $\Phi_t^i(\cdot)$ contains a separate model for the controllable and uncontrollable demand, i.e., $\Phi_t^i(\cdot) = \{\Phi_t^{\mathrm{c},i}(\cdot), \Phi_t^{\mathrm{uc},i}(\cdot)\}$. The modified prediction $\widetilde{\theta}_t^i$ is an adjusted version of $\widehat{\theta}_t^i$ that incorporates information from the newly arrived measurement y_t using a gradient descent-based update. The specified model $\Phi_t^i(\cdot)$ advances the estimate $\widetilde{\theta}_t^i$ in time. The gradient descent-based then model-based update process is analogous to a Kalman filter's separate steps of incorporating a new observation then advancing the prediction using a model. In contrast with a Kalman filter, the models can have an arbitrary structure and there are no assumptions on the underlying noise distribution.

DMD forms predictions for models $i \in \mathcal{N}^{\mathrm{mdl}}$ as

$$\widetilde{\theta}_t^i = \underset{\theta\in\Theta}{\arg\min}\ \eta_t \left\langle \nabla \ell_t(\widehat{\theta}_t^i, y_t),\ \theta \right\rangle + D\left(\theta \| \widehat{\theta}_t^i\right), \quad (6)$$

$$\widehat{\theta}_{t+1}^i = \Phi_t^i(\widetilde{\theta}_t^i), \quad (7)$$

where (6) incorporates the measurement into the estimate and (7) advances the estimate in time using the model. We

minimize over the variable θ. The term $\eta_t\langle\nabla\ell_t(\widehat{\theta}_t^i, y_t),\ \theta\rangle$ captures the alignment of the variable θ with the positive gradient of a convex loss function $\ell(\widehat{\theta}_t^i, y_t)$. To minimize this term alone, we would choose θ to be exactly aligned with the negative gradient direction. The term $D(\theta\|\widehat{\theta}_t^i)$ is a Bregman divergence that penalizes the deviation between the new variable θ and the old variable $\widehat{\theta}_t^i$. It can be thought of as a term that trades-off matching the noisy data points with trusting the predictions. The value η_t is a step-size parameter and $\langle\cdot,\cdot\rangle$ is a dot product. Within this work, we set $\ell_t(\widehat{\theta}_t^i, y_t) = 10^{-4}\cdot\|y_t - (\widehat{y}_t^{c,i} + \widehat{y}_t^{uc,i})\|_2^2$ and $D(\theta\|\widehat{\theta}_t^i) = \|\theta - \widehat{\theta}_t^i\|_2^2$ where we include 10^{-4} in the loss function for numerical reasons.

The second DMD process forms the overall estimate from a weighted combination of the N^{mdl} predictions

$$\widehat{\theta}_{t+1} = \sum_{i=1}^{N^{\text{mdl}}} w_{t+1}^i \widehat{\theta}_{t+1}^i. \tag{8}$$

The weights are based on each model's historical accuracy with respect to the measurements y_t, and models that incur larger loss values have less influence in the overall estimate. The algorithm generating the weights is

$$w_{t+1}^i = \frac{\lambda}{N^{\text{mdl}}} + (1-\lambda)\widetilde{w}_t^i \tag{9}$$

$$\widetilde{w}_t^i = \frac{w_t^i\ \exp\left(-\eta^r\,\ell_t\left(\widehat{\theta}_t^i, y_t\right)\right)}{\sum_{j=1}^{N^{\text{mdl}}} w_t^j\ \exp\left(-\eta^r\,\ell_t\left(\widehat{\theta}_t^j, y_t\right)\right)} \tag{10}$$

where w_t^i is the weight associated with model i at time-step t, $\lambda \in (0,1)$ dictates the portion of weighting that is shared among models, and $\widetilde{w}_t^i$ is the preliminary weight for model i based on the loss of each model and the total loss of all models.

Ref. [7] describes λ as a parameter that allows fast switching between models. Another aspect of this is that λ sets the default combination of models used for predictions. For example, when λ is nearly one, all N^{mdl} models have almost equal share in the overall estimate regardless of their previous loss. We investigated DMD's accuracy as we increase λ, and the RMS error in the controllable demand estimate increases with larger values for λ. Understanding this behavior in more detail and investigating an alternative weighting function that avoids this issue are proposed for future work.

V. MODELS USED WITHIN DMD

Each of the N^{mdl} models used within the DMD algorithm consists of a controllable demand model paired with an uncontrollable demand model. We assume the air conditioner measurements transmitted to the aggregator, e.g., at the end of the day, by a house's smart meter are a history of the house's internal air temperature and the air conditioner's on/off mode $y_t^{\text{ac},i} = \begin{bmatrix}\theta^{a,i} & m_t^i\end{bmatrix}^T$. We do not include the house's internal mass temperature as it is not available for measurement. We use this historical data to build the various controllable and uncontrollable load models used within the algorithm, and we detail these models in this section.

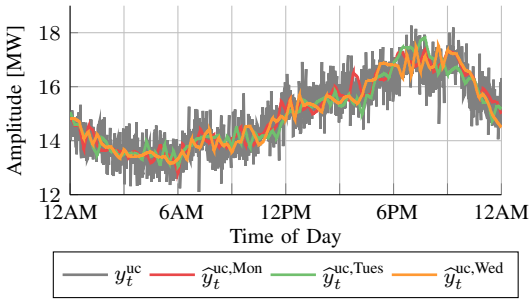

Fig. 3. The plant's true uncontrollable demand and three predictions of the uncontrollable demand.

A. Uncontrollable Demand Model

There exist many methods of forecasting demand, for example [20], which takes the time of week and weather into account. While advanced models may be accurate more frequently, they will invariably be incorrect at some point. DMD has the ability to find an alternative model in these scenarios, provided an alternative model or combination of models is accurate. For this reason we include simpler forecasting methods and rely on DMD to select and switch between the most accurate model (or combination of models) in real-time.

The uncontrollable demand model is simply a lookup table

$$\widehat{y}_t^{uc} = \alpha_t. \tag{11}$$

The α_t values are predetermined power values for each time-step t over the desired prediction horizon. Within this work, an α_t value exists for each one minute interval over the course of the day, and we generate the α_t values from some previous day's true uncontrollable demand. The true uncontrollable demand is found by removing the total controllable demand, which is known from the transmitted air conditioner data, from the total demand measurements. After constructing the uncontrollable demand signal for a previous day, the time series is broken into fifteen minute intervals. A linear least squares fit is generated for each fifteen minute interval, and we combine the linear predictors into a piecewise linear, continuous function. The α_t values are this function's predicted uncontrollable demand for each time of day.

We generate the uncontrollable demand models using uncontrollable demand data from the Dataport for July 27-29 – the Monday, Tuesday, and Wednesday of the week preceding the simulated day. We denote their uncontrollable demand predictions as $\widehat{y}_t^{\text{uc,Mon}}$, $\widehat{y}_t^{\text{uc,Tues}}$, and $\widehat{y}_t^{\text{uc,Wed}}$ respectively, and the models are denoted $\Phi_t^{\text{uc,Mon}}$, $\Phi_t^{\text{uc,Tues}}$, and $\Phi_t^{\text{uc,Wed}}$ respectively. Figure 3 depicts the plant uncontrollable demand y_t^{uc} and the three uncontrollable demand predictions.

B. Controllable Demand Models

We use three controllable demand models within the DMD algorithm. Similar to the construction of the plant's controllable demand signal, the first uncontrollable demand model is a collection of hybrid models where the individual hybrid models capture individual air conditioners' power

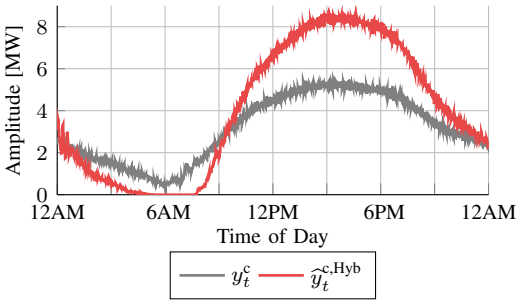

Fig. 4. Time series plots comparing of the plant's controllable demand to the predicted demand of the two-state hybrid model population.

draw. In contrast to the hybrid models within the plant, these hybrid models have less detail. The predicted controllable demand at a given time is the sum of each hybrid model's demand.

The second two models are versions of an aggregate model developed in [11]. Whereas each hybrid model attempts to capture the heat transfer driving the air conditioner's cycling, the aggregate model seeks to describe the dynamics of the entire air conditioner population probabilistically. The first aggregate model is a set of linear time-invariant (LTI) models that characterize air conditioner population dynamics based on different outdoor temperatures. The second aggregate model is a linear time-varying (LTV) model generated from a previous day's air conditioner data.

Note that all of the controllable demand models operate with two second time-steps, and each model's predictions within a one minute interval are averaged to form the corresponding controllable demand predictions. In addition, we initialize each model with the "true", i.e., plant, values. The following subsections detail the controllable demand models.

1) Hybrid Model Population: Since we do not have access to internal mass temperature measurements, the hybrid model states include only the house's internal temperature, i.e., $\theta_t^i = \theta_t^{\mathrm{a},i}$ and the on/off mode of the appliance m_t^i. The vector of environmental heat sources is the outdoor temperature $d_t^i = \theta_t^{\mathrm{o}}$.

The resulting continuous-time matrices are

$$A^{\mathrm{c},i} = -\widehat{U}^{\mathrm{a},i}/\widehat{\Lambda}^{\mathrm{a},i} \tag{12}$$

$$B^{\mathrm{c},i} = Q^{\mathrm{h},i}/\widehat{\Lambda}^{\mathrm{a},i} \tag{13}$$

$$E^{\mathrm{c},i} = \widehat{U}^{\mathrm{a},i}/\widehat{\Lambda}^{\mathrm{a},i} \tag{14}$$

where $\widehat{\Lambda}$ and $\widehat{U}$ are identified using a nonlinear least squares algorithm. The nonlinear least squares algorithm inputs are 1) historical data for the house's internal air temperature, 2) historical data for the air conditioner's on/off mode, and 3) the corresponding outdoor air temperature data for a given appliance or household. The predicted power draw for an individual model is $\widehat{P}_t^i = (|Q^{\mathrm{h},i}|\ m_t^i)/\eta^i$ where $Q^{\mathrm{h},i}$ and η^i are assumed to be known.

The predicted controllable demand at each time-step is $\widehat{y}_t^{\mathrm{c,Hyb}} = \sum_{i=1}^{N^{\mathrm{ac}}} \widehat{P}_t^i$. Because the hybrid model population results in a single controllable demand prediction, we note it collectively as $\Phi_t^{\mathrm{c,Hyb}}$.

Figure 4 shows the modeled controllable demand versus the plant's controllable demand using $\widehat{\Lambda}$ and $\widehat{U}$ values that are identified once at the beginning of the day. Note that the model performs poorly over the course of the day; its resulting weight in the DMD algorithm is small. Future work should investigate estimating the environmental heat injections, which would allow a more detailed d_t^i vector in the two-state model. Additional future work should investigate hybrid system identification techniques that allow identification of the model parameters and model order based on input-output data.

2) LTI Aggregate Model Set: The aggregate model for a controllable appliance population predicts the portion of appliances that draw power versus the portion of appliances that are not drawing power. It does so by forming a state vector $x \in \mathbb{R}^{N^{\mathrm{x}}}$ that consists of the portion of controllable appliances within discrete state bins. In this work, one bin captures the portion of controllable appliances that are on and another captures those that are off, i.e., $N^{\mathrm{x}} = 2$. The state transition matrix, $A \in \mathbb{R}^{N^{\mathrm{x}} \times N^{\mathrm{x}}}$, is a transposed Markov transition matrix. Its entries capture 1) the probability that appliances stay within a given state bin during the time-step, and 2) the probability of switching state bins. The output matrix is $C = N^{\mathrm{ac}} \overline{P} \begin{bmatrix} 0 & 1 \end{bmatrix}$ where $\overline{P}$ is the average power draw of appliances that are on.

We denote a set corresponding to outdoor temperatures of interest as $\mathcal{J}^{\mathrm{temps}}$. The set consists of integer temperatures between an upper temperature of interest $\overline{\theta}^{\mathrm{o}}$ and a lower temperature of interest $\underline{\theta}^{\mathrm{o}}$. The $\underline{\theta}^{\mathrm{o}}$ value is selected lower than the minimum temperature of the simulated day, and the $\overline{\theta}^{\mathrm{o}}$ is selected greater than the maximum temperature.

The LTI model set is generated by constructing A^j and C^j matrices for $j \in \mathcal{J}^{\mathrm{temps}}$ using air conditioner on/off data at each temperature to identify the corresponding matrices. The output $\widehat{y}_t^{\mathrm{c,LTI},j}$ is the predicted power draw of the controllable appliance population. The resulting model is

$$x_{t+1}^j = A^j\, x_t^j \qquad j \in \mathcal{J}^{\mathrm{temps}} \tag{15}$$

$$\widehat{y}_t^{\mathrm{c,LTI},j} = C^j\, x_t^j \qquad j \in \mathcal{J}^{\mathrm{temps}}. \tag{16}$$

Note that the state of each model evolves independently of the states of the other models. We note the set of LTI models as $\Phi_t^{\mathrm{c,LTI}}$and their predictions as $\widehat{y}_t^{\mathrm{c,LTI}}$.

3) LTV Aggregate Model: This approach uses a previous day's controllable appliance data to develop a trajectory of A_t matrices. We associate an A_t matrix with each time-step, and we generate the matrix using the controllable appliances' historical on/off transitions within that time-step. Note that the state transition matrix implicitly has a dependence on the outdoor air temperature; hotter temperatures will force more appliances to draw power at a given time.

The output matrix is $C_t = N^{\mathrm{ac}} \overline{P}_t \begin{bmatrix} 0 & 1 \end{bmatrix}$ where $\overline{P}_t$ corresponds to the average demand of appliances that are on during a given time-step. The corresponding model is

$$x_{t+1} = A_t\, x_t \tag{17}$$

$$\widehat{y}_t^{\mathrm{c,LTV}} = C_t\, x_t. \tag{18}$$

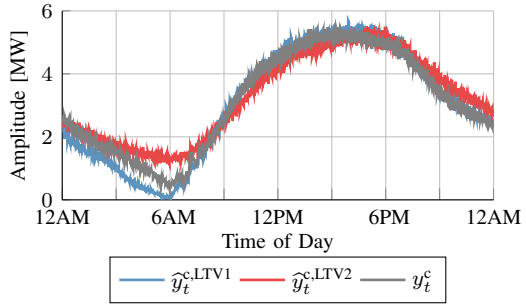

Fig. 5. Time series plots comparing of the actual controllable demand sampled to two LTV aggregate models generated from data corresponding to different outdoor temperature trajectories.

The historical state transition matrix trajectories will in general not be accurate unless the current day's outdoor temperature follows the same trajectory as the one used to identify the A_t matrices. We develop two aggregate LTV models, denoted $\Phi_t^{\mathrm{c,LTV1}}$ and $\Phi_t^{\mathrm{c,LTV2}}$, from the outdoor temperature trajectories in Fig. 2. Figure 5 depicts their predictions, $\widehat{y}_t^{\mathrm{c,LTV1}}$ and $\widehat{y}_t^{\mathrm{c,LTV2}}$, respectively. Future work should formulate a state transition matrix that is an explicit function of the outdoor air temperature.

VI. CASE STUDIES

The following subsections provide several case studies that explore the effectiveness of DMD within our scenario. First, we present a benchmark case that uses a simple interpolation-based algorithm, not DMD, to generate controllable demand predictions based on the outdoor temperature. Following this, we present a series of simulations that illustrate DMD's capabilities and limitations in the given scenario. We then explore how performance changes when varying η^{r}.

In all cases, the controllable demand model states are initialized at the true, i.e., plant, value, and the models run open-loop, i.e., without real-time measurements. We use the RMS error associated with the controllable demand estimate to quantify the performance, and Table II summarizes the cases and results.

A. Benchmark Case

In this section, we provide the results for a simple algorithm that estimates the controllable demand based on the outdoor air temperature and a collection of aggregate LTI models. These results serve as a baseline for comparison with the DMD algorithm in the following subsections.

The simple algorithm uses the steady-state controllable demand y_{ss}^j, predicted by the LTI aggregate models $j \in \mathcal{J}^{\mathrm{temps}}$. Given an outdoor temperature measurement θ_t^{o}, the controllable demand is estimated via interpolation as

$$\widehat{y}_t^{\mathrm{c}} = y_{\mathrm{ss}}^- + (y_{\mathrm{ss}}^+ - y_{\mathrm{ss}}^-)\frac{\theta_t^{\mathrm{o}} - \theta^{\mathrm{o},-}}{\theta^{\mathrm{o},+} - \theta^{\mathrm{o},-}} \tag{19}$$

where $\theta^{\mathrm{o},+} = \lceil \theta_t^{\mathrm{o}} \rceil$ and $\theta^{\mathrm{o},-} = \lfloor \theta_t^{\mathrm{o}} \rfloor$. The $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$ operators round up and down respectively. The y_{ss}^+ and y_{ss}^- values are the steady-state demand of the aggregate LTI model corresponding to $\theta^{\mathrm{o},+}$ and $\theta^{\mathrm{o},-}$ respectively. This method estimates the controllable demand with an RMS error of 738 kW, and the time series plots are not included.

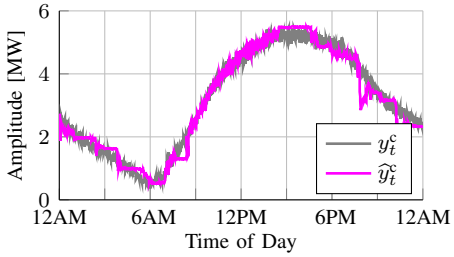

Fig. 6. Controllable demand estimate from the DMD Case 1.

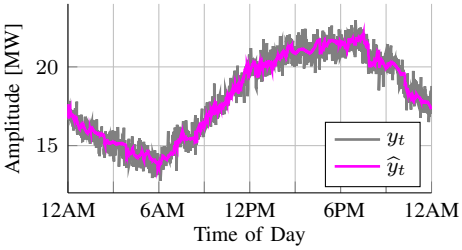

Fig. 7. Total demand estimate from DMD Case 1.

B. DMD Cases

In this section we present several simulation cases in which we vary the models used to illustrate properties of DMD. First, we present an example of DMD's effectiveness in estimating the controllable demand using a realistic set of models. Following this, we introduce additional, unrealistic models that show how model accuracy can influence the estimation results. Finally, we present a cautionary case that illustrates a potential problem with DMD in this scenario.

1) DMD Case 1: The set of N^{mdl} models used in this case consists of every possible controllable and uncontrollable demand model pair where the models are listed in Table II. Additional parameter settings include $\lambda = 0.001$, $\eta^r = 0.03$, $\eta_t = 0.34$,

Figures 6 and 7 provide time series showing the controllable demand estimate and total demand estimate, respectively. Each figure includes the corresponding plant signals. The RMS error of the controllable demand estimate is 264 kW. The accuracy is noteworthy in that none of the controllable load models included in the case are accurate between 12 and 8 AM, but the algorithm still does fairly well in estimating the controllable demand over this time period. The step-like shapes in Fig. 6, e.g., at 3-6 AM, are due to DMD relying on and transitioning between LTI models. Generating models for smaller temperature intervals improves this accuracy, which leads us to believe an aggregate A matrix that is a function of outdoor temperature should improve the estimation performance. Finally, the estimation errors from from 3 PM to 12 AM are mainly attributable to uncontrollable demand modeling inaccuracies.

2) DMD Case 2: Including models that are more accurate, or produce more accurate models combinations, can help reduce the RMS error. However, the additional models are only useful if they provide additional accuracy during times

TABLE II
SUMMARY OF RESULTS

Simulation Name	Section	Algorithm	Models	RMS Error [kW]
Benchmark Case	VI-A	Interpolation	$\Phi_t^{c,LTI}$	738
DMD Case 1	VI-B.1	DMD	$\Phi_t^{c,LTI}$, $\Phi_t^{c,LTV1}$, $\Phi_t^{c,LTV2}$, $\Phi_t^{c,Hyb}$, $\Phi_t^{uc,Mon}$, $\Phi_t^{uc,Tues}$, $\Phi_t^{uc,Wed}$	264
DMD Case 2	VI-B.2	DMD	DMD Case 1 Models, $\Phi_t^{uc,Today}$	211
DMD Case 3	VI-B.3	DMD	DMD Case 2 Models, $\Phi^{c,PA}$	175
DMD Case 4	VI-B.4	DMD	$\Phi_t^{uc,UE}$, $\Phi_t^{c,LTI}$, $\Phi_t^{c,LTV1}$, $\Phi_t^{c,LTV2}$, $\Phi_t^{c,Hyb}$	1392
η^r Sweep	VI-C	DMD	DMD Case 1 Models	264-343

when the existing models were inaccurate.

For example, we generate a new uncontrollable demand model by applying the method in Section V-A to the simulated day's actual uncontrollable demand. We denote the model as $\Phi_t^{uc,Today}$, and this model will be accurate over the entire day due to its construction. The models available within DMD in this case include those within DMD Case 1 plus $\Phi_t^{uc,Today}$, and the model combinations used consists of every possible pair of these controllable and uncontrollable models. The RMS error is reduced to 211 kW, and including an accurate uncontrollable demand model fixes the majority of the biases between 3 PM and 12 AM in Fig. 6.

3) DMD Case 3: This simulation illustrates the point that including additional models whose accuracy complements that of the existing models improves the overall estimation performance. Recall that the controllable demand models are inaccurate over the first eight hours of the day. To complement this, we generate a new controllable demand model whose predicted demand values for the first 8 hours of the day correspond exactly to the plant's controllable demand. The model's controllable demand predictions for the final 16 hours are 0 kW. Clearly, this model is accurate over the first 8 hours where previous controllable demand models were inaccurate, and it is inaccurate over the remaining 16 hours where previous models were accurate.

We denote this model as $\Phi^{c,PA}$ and include it in the set of models for DMD Case 2. DMD's model set consists of all possible pairs of controllable and uncontrollable models within this set. Including a controllable demand model that is accurate where the others are inaccurate reduces the estimation error further to 175 kW by fixing many of the biases between 12 and 8 AM.

Figure 8 shows the evolution of some of the model weights in this case. There are too many models within the case to include explicitly, and so we plot the few heavily utilized models along with the "Other Models" time series that consists of the sum of all other weights. We label the weights using the controllable demand model component, and the corresponding uncontrollable demand model component of each weight is $\Phi_t^{uc,Today}$. Predictions $\widehat{y}_t^{c,LTI,36}$ correspond to the LTI model for 36°C. Note that the dominant, i.e., most heavily weighted model, switches as the scenario progresses. Also note that there are times when the "Other Models" time series is dominant; this consists of times when a combination of a number of models are used to form the overall estimate.

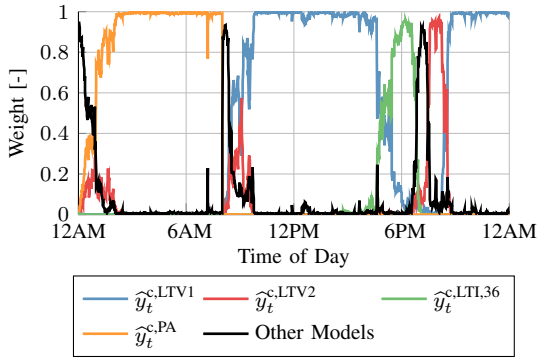

Fig. 8. Example evolution of select model weights in DMD Case 3. The uncontrollable model component of each weight is $\Phi_t^{uc,Today}$, and $\widehat{y}_t^{c,LTI,36}$ indicates predictions of the LTI model for 36°C

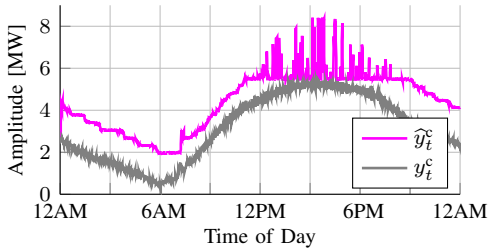

Fig. 9. Example from DMD Case 4 of DMD compensating if all uncontrollable model predictions are too low.

4) DMD Case 4: In the final case, we show a limitation of DMD in our scenario. Specifically, because the observations and model outputs all lie within $\mathbb{R}$, the algorithm simply tries to find the model combination whose total demand prediction closely matches the observations. If one component model underestimates the demand while the other overestimates the demand, the total demand predictions can still closely match the measurements and the model can be weighted heavily. To illustrate this, we introduce an uncontrollable model $\Phi_t^{uc,UE}$ whose estimates are $\widehat{y}_t^{uc,UE} = 0.9 \cdot \widehat{y}_t^{uc,Today}$, and we do not include the other uncontrollable demand models in the algorithm.

Figure 9 shows how DMD overestimates the controllable demand to compensate, and the resulting RMS error in the controllable demand estimate is 1392 kW. Note that between 12 and 8 PM the algorithm switches between relying on the hybrid model population's controllable demand prediction and an LTI model's demand prediction. Considering apparent power, which can be represented in $\mathbb{R}^2$, should add additional robustness against this issue, and investigating this extension is the topic of future work.

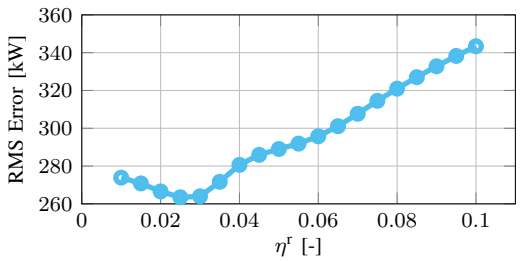

Fig. 10. Plot from the η^r sweep depicting the variation in the RMS estimation error of the controllable demand.

C. Investigating the Effect of η^r

The η^r parameter appears in (10), which updates the model weights. In this section, we describe the impact of varying η^r on the estimate accuracy. The models correspond to that in DMD Case 1. Figure 10 shows how the estimated controllable demand's RMS error varies with the parameter's value.

Larger values of η^r tend to allow faster switching between models for a given λ value. This makes intuitive sense as it appears as a negative power of e within (10). However, as this value increases, the weights can vary rapidly, and DMD attempts to match the noisy data too closely. The corresponding estimated demand appears more noisy.

Decreasing the value of η^r slows down the transitions between models. This is advantageous as it allows a smoother evolution of the weights and the corresponding estimates. However, it can be restrictive in that the switching between the dominant model combination may be too slow. This trade-off dictates the shape of the plot in Fig. 10.

VII. CONCLUSIONS

In this paper, we introduced the application of online learning algorithms to deduce DERs behavior from real-time, feeder-level measurements. Specifically, we explored using DMD to deduce the active power demand of a population of controllable appliances from feeder-level demand measurements that include uncontrollable demand.

We developed a simulation environment using data from GridLAB-D, data from the Pecan Street Inc. Dataport, and models for controllable appliances. We then presented DMD – the online learning algorithm – and we presented several models that were used within DMD. The results show that DMD can effectively determine the portion of total demand measurements that correspond to controllable demand. The DMD algorithm outperformed an interpolation-based algorithm that relies on a set of LTI aggregate models, and incorporating more accurate models into DMD improves the estimate of the controllable demand, even if those models are not accurate over the entire day. Finally, we showed a shortcoming of DMD within the proposed scenario.

Addressing this shortcoming, improving controllable load models, and modifying the DMD algorithm are several avenues of future research. Specifically, extending the approach to handle apparent power measurements, rather than active power measurements, should help make the DMD algorithm more robust in our scenario. Future work in modeling includes 1) developing an outdoor-temperature-dependent state transition matrix in the aggregate model, 2) introducing more realistic environmental heat injections into the individual appliance models, and 3) applying more general system identification methods to determine better models for individual household air conditioners based on the readily available data. Finally, future work should investigate alternative weighting functions for the DMD algorithm.

REFERENCES

[1] GTM Research/SEIA: U.S. Solar Market Insight, "Solar market insight report 2015 Q1," http://www.seia.org/research-resources/solar-market-insight-report-2015-q1, 2015.

[2] Navigant Research, "Direct load control and dynamic pricing programs, DR markets, and DR management systems for residential customers: Global market analysis and forecasts," https://www.navigantresearch.com/research/residential-demand-response#, 2015.

[3] K. Armel, A. Gupta, G. Shrimali, and A. Albert, "Is disaggregation the holy grail of energy efficiency? The case of electricity," *Energy Policy*, vol. 52, pp. 213–234, 2013.

[4] R. Dong, L. Ratliff, H. Ohlsson, and S. Sastry, "Fundamental limits of nonintrusive load monitoring," in *Proceedings of the HiCoNS*, 2014.

[5] J. Kolter and T. Jaakkola, "Approximate inference in additive factorial HMMs with application to energy disaggregation," in *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 2012.

[6] R. Dong, L. Ratliff, H. Ohlsson, and S. Sastry, "Energy disaggregation via adaptive filtering," in *Proceedings of the Allerton Conference*, 2013.

[7] E. Hall and R. Willett, "Online convex optimization in dynamic environments," *IEEE Journal of Selected Topics in Signal Processing – Signal Processing for Big Data*, 2015, to appear.

[8] J. Mathieu, M. Dyson, and D. Callaway, "Resource and revenue potential of California residential load participation in ancillary services," *Energy Policy*, vol. 80, pp. 76–87, 2015.

[9] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent, "Aggregate flexibility of thermostatically controlled loads," *IEEE Transactions on Power Systems*, vol. 30, no. 1, pp. 189–198, 2015.

[10] S. Koch, J. Mathieu, and D. Callaway, "Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services," in *Proceedings of the Power Systems Computation Conference*, Stockholm, Sweden, Aug. 2011.

[11] J. Mathieu, S. Koch, and D. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Transactions on Power Systems*, vol. 28, no. 1, pp. 430–440, 2013.

[12] E. Can Kara, Z. Kolter, M. Berges, B. Krogh, G. Hug, and T. Yuksel, "A moving horizon state estimator in the control of thermostatically controlled loads for demand response," in *Proceedings of SmartGridComm*, Vancouver, BC, 2013.

[13] E. Vrettos, J. Mathieu, and G. Andersson, "Control of thermostatic loads using moving horizon estimation of individual load states," in *Proceedings of the Power Systems Computation Conference (PSCC)*, Wroclaw, Poland, Aug. 2014.

[14] G. S. Ledva, E. Vrettos, S. Mastellone, G. Andersson, and J. L. Mathieu, "Applying networked estimation and control algorithms to address communication bandwidth limitations and latencies in demand response," in *Hawaii International Conference on Systems Science (HICSS)*, Grand Hyatt, Kauai, HI, 2015.

[15] "GridLAB-D Feeder Taxonomy Documentation," http://gridlab-d.sourceforge.net/wiki/index.php/Feeder_Taxonomy, 2015.

[16] Pecan Street Inc., "Dataport," https://dataport.pecanstreet.org/, 2015.

[17] R. C. Sonderegger, "Dynamic models of house heating based on equivalent thermal parameters," Ph.D. dissertation, Princeton Univ., NJ., 1978.

[18] "GridLAB-D House Class Documentation," http://gridlab-d.sourceforge.net/wiki/index.php/House, 2015.

[19] K. Ogata, *Discrete-time control systems*. Prentice Hall Englewood Cliffs, NJ, 1995, vol. 2.

[20] J. L. Mathieu, P. N. Price, S. Kiliccote, and M. A. Piette, "Quantifying changes in building electricity use, with application to demand response," *IEEE Transactions on Smart Grid*, vol. 2, no. 3, pp. 507–518, 2011.