

A Sparse Linear Model and Significance Test for Individual Consumption Prediction

Pan Li, Baosen Zhang, Yang Weng, and Ram Rajagopal

Abstract—Accurate prediction of user consumption is a key part not only in understanding consumer flexibility and behavior patterns, but in the design of robust and efficient energy saving programs as well. Existing prediction methods usually have high relative errors that can be larger than 30% and have difficulties accounting for heterogeneity between individual users. In this paper, we propose a method to improve prediction accuracy of individual users by adaptively exploring sparsity in historical data and leveraging predictive relationship between different users. Sparsity is captured by popular least absolute shrinkage and selection estimator, while user selection is formulated as an optimal hypothesis testing problem and solved via a covariance test. Using real world data from PG&E, we provide extensive simulation validation of the proposed method against well-known techniques such as support vector machine, principle component analysis combined with linear regression, and random forest. The results demonstrate that our proposed methods are operationally efficient because of linear nature, and achieve optimal prediction performance.

Index Terms—Load forecasting, least absolute shrinkage and selection, sparse autoregressive model, significance test

I. INTRODUCTION

ELECTRIC load forecasting is an important problem in the power engineering industry and have received extensive attention from both industry and academia over the last century. Many different forecasting techniques have been developed during this time. The authors in [1] present a comprehensive literature review on different methods related to load forecasting, from regression models to expert systems. Time series methods are further discussed in [2]. A thorough research on load and price forecasting is presented in [3]. A common theme among many of the established methods is that they are used to forecast relative large loads, from substations serving megawatts to transmission networks serving more than gigawatts of power [4]. As the grid shifts to a more distributed system, the need of accurate forecasting for smaller sized loads is becoming increasing important.

Recent advances in technology such as smart meters, bi-directional communication capabilities and distributed energy resources have made individual households active participants in the power system. Many applications and programs based on these new technologies require estimating the future load of individual homes. For example, state estimation algorithms for distribution systems require pseudo-measurements [5],

Pan Li and Baosen Zhang are with the Department of Electrical Engineering, University of Washington, Seattle, WA, 98195, (e-mail: {pli69, zhangba}@uw.edu).

Yang Weng and Ram Rajagopal are with the Civil and Environmental Department, Stanford University, Stanford, CA, 94035, (e-mail: {yangweng, ramr}@stanford.edu).

and these are provided by load forecasts. Another important class of application is demand response and dynamic pricing programs, where users' future demand are needed to design appropriate incentives [6], [7], [8], [9], [10].

In contrast to large aggregated loads, individual load forecasting is less developed. The current state of forecasting algorithms falls under three broad classes: simple averaging, statistical regression methods, and Artificial Intelligence (AI). They are listed in increasing order of prediction accuracy and decreasing in order of model simplicity. Simple averaging is intuitively pleasing since it is based on the mean of the previous similar days, but is often not very accurate. On the other end of the spectrum, AI methods can be extremely accurate, but it is difficult to associate the obtained parameter with the input data, i.e., past consumption. In this paper, we propose an algorithm that can achieve the performance of the state of the art AI methods, but retains the simplicity of linear nature of regression methods.

The algorithm we propose is based on the well known *Least Absolute Shrinkage and Selection (LASSO)* algorithm in statistics and signal processing [11]. It shrinks parameter estimates towards zero in order to avoid overfitting as well as pick up the most relevant regressors. Operationally, it minimizes the usual sum of squared errors, with a bound on the sum of the absolute values of the coefficients. LASSO is preferable in a setting where the dimension of the features are much higher than the size of the training set [12], [13]. Furthermore, it can be easily extended to suit for more scenarios, i.e., group LASSO or adaptive LASSO [14], [15]. Due to its nice statistical properties and its efficiency, LASSO is applied to many different disciplines and has shown its superiority in sparsity recovery [16], [17].

In particular, we formulate the load forecasting problem as learning parameters of a *sparse* autoregressive model using LASSO. This sparse autoregressive model automatically selects the best recurrent pattern in historical data by shrinking irrelevant coefficients to zero. By selecting the correct features, the algorithm improves the order selection in autoregressive models, and as we will show using real load data from Pacific Gas and Electric Company (PG&E), it also greatly improves the prediction accuracy of current regression models. For example, taking the Mean Absolute Percentage Error (MAPE) as a metric, autoregressive model with lag order one (AR(1)) has a MAPE of 33.9 %, based on a pool of 150 users. The proposed method reduces the MAPE to 22.5 %. Nonparametric methods such as using the consumption data at the same time during last week (LW method) are more intuitive but their performance are not robust subject to noise. They yield

a relative error nearly up to 35% based on a pool of 50 users in our test dataset. In addition, based on our dataset, linear regression model with the reduced feature space from Principle Component Analysis (PCA) performs rather poorly, with a relative error greater than 100 %, which is even worse than simply predicting zero. Overall, the proposed model with LASSO achieves a much smaller error and maintains the intuitive nature of linear methods.

In addition, we observe that given the set of coefficients, utilities and planners can easily identify and attribute the impact of different features on load consumption based on LASSO. For example, suppose for a given user, our method concludes that its load at time t is mostly determined by the past load at times $t - 1$, $t - 3$ and $t - 24$. This shows that this user has some short term behavior in the one hour range, medium term behavior in the 3 hour range, and a daily cycle partner that repeats every 24 hours. This information can be interpreted by utilities to better understand this user's consumption patterns and potentially identify appropriate demand programs.

Apart from LASSO, data from other users can be leveraged for forecasting in the proposed framework. Intuitively, this means that knowing the past history of user j improves forecasting of the user i . For a given user, we use a sequential hypothesis testing procedure to find the best other user's historical data to include in the algorithm. We give a rigorous derivation of the hypothesis testing procedure and quantify the confidence of including other users. This allows us to show that the procedure is optimal, in the sense that given user i , it finds user j , whose historical data improves load forecasting for user i the most among all users.

We derive rigorous theoretical justification for our methods as well as provide extensive simulation studies with respect to several well studied prediction methods. In particular, we compare against AutoRegression (AR), Exponential Smoothing (ES), Support Vector Machine (SVM), linear regression with PCA, Random Forest (RF) and Neural Network (NN) model. We also include two more nonparametric methods, i.e., using ten previous days' average consumption as prediction and using last week's (LW as mentioned earlier) consumption data as prediction. Using a user's own historical data, our proposed method and RF both reduce prediction error by 30% compared to other predictors. Our proposed method is simpler in nature than RF since the latter is a generic machine learning technique that relies on a random ensemble of decision trees [18], [19]. Therefore, our method is useful to system operators in policy decisions without sacrificing prediction accuracy. By adding the historical data of another user, we can further improve the prediction accuracy.

A. Contribution

The contribution of this paper is two fold.

- First, we apply LASSO in consumption prediction. LASSO has been studied in literature to jointly predicts price, load and renewables together [20]. However, there is little discussion on how LASSO can be applied to consumption forecasting without using other side information. Our paper has a thorough discussion on how

LASSO is used in obtaining a sparse autoregressive model and compares its performances with several other popular prediction methods in literature. The simulation results show that LASSO achieves competitive prediction performance compared to nonlinear machine learning algorithms such as RF or NN. In addition, LASSO is computationally much faster and easier to be understood by human operators.

- Second, we propose a significance test that can leverage other users' consumption data for prediction. This testing procedure differs significantly from standard clustering algorithms since it looks for the most "predictive" user, not necessarily the most similar user.

The rest of the paper is organized as follows. Section I-B analyzes related work in short term load forecasting. Section II presents the autoregressive model for time series analysis. Section III introduces LASSO type linear regression model. Section IV proceeds with the significance test to pair users in order to improve prediction performance. It describes the significance test for LASSO, i.e., covariance test, to select the most significant user to form the pair with the current user. Section V introduces the evaluation criterion for prediction and details the simulation of the proposed methods compared to several other popular prediction methods. And finally Section VI concludes the paper and draws avenues for future work.

B. Literature review

There exists an extensive literature on short term load forecasting [21], [1]. In summary, research on energy consumption prediction can be divided into three groups [1], including simple averaging models, statistical models and AI models.

The simplest approach is to employ moving average [22]. Such models make predictions on mean of consumption data from previous similar days [23]. AI type methods (e.g. NN or RF) yield high accuracy at the cost of complexity of the system, which may lead to overfitting [24]. Other drawbacks include difficult parametrization and non-obvious selection of variables. Statistical methods sit in between the previous methods in terms of complexity and accuracy, and include regression models, probabilistic approach applied to regression models, and time series analysis such as autoregressive models.

In statistical methods, regression models combine several features to form normally a linear function. In [25], the authors build a regression tree model with weather data to predict consumption. SVM is used in [26]. Gaussian process framework for prediction mitigating the uncertainty problem is proposed recently in [27].

Advanced AI models have also been used to facilitate load forecasting. Authors in [28] propose a NN model to handle uncertainties associated with forecasts. Similarly, authors in [29] study the interruption prediction with weather conditions by simulating a NN. Besides these work, the authors in [30] explore kernel-based learning techniques to forecast both commercial and industrial load. In addition, feature selection techniques are applied in [31] to predict electricity load and prices.

Apart from these advanced machine learning tools, another type model involves time series based methods [32]. An overview can be found in [33]. Models such as ES [34] and the Autoregressive Integrated Moving Average (ARIMA) model [35], [36] are all time series based. ARIMA is a widely used time series based prediction methods and have been adopted in [37], [38]. In these papers the authors apply ARIMA to predict either the electricity price or the electric vehicle demand. Apart from these work, authors in [20] propose a vector autoregressive model to include renewables, prices and loads together with sparsity recovery. In this work, the authors have also explored LASSO to obtain a sparse linear model. In addition, to extend from linearity to nonlinearity, [39] addresses a mixed model combining ARIMA model to deal with the linear part and neural network with the nonlinear one. In our work, we recover the sparsity for univariate time series and multivariate time series under the framework of autoregressive models.

II. AUTOREGRESSIVE MODEL

Autoregressive models are widely used for prediction and inference of time series data. Here we adopt a linear autoregressive model of the hourly consumption of a single user, where future demands y_t depend linearly on historical data y_{t-i} plus random noise:

$$y_t = \beta_0 + \sum_{i=1}^I \beta_i y_{t-i} + \varepsilon_t. \quad (1)$$

In this model, y_t denotes the demand of the user at time t , β_i is the coefficient for order (lag order) i in the autoregressive model (it represents the weight of each historical demand data y_{t-i} in predicting future demand y_t), and ε_t is an additive random Gaussian noise. The time index t is measured in hours and the noise is identically and independently distributed at different hours. Note that in this paper we denote time series data by notation $\{\bullet_t\}$, where subscript t refers to the time slot in this time series data. In addition, I is the number of orders that we include in the model. An autoregressive model with maximum order I is denoted by AR(I).

To use the model in (1) for prediction, the standard approach is to use Ordinary Least Squares (OLS) to estimate the coefficients β_i , $i \in I$. By convention, we write $[y_{t-1} \ y_{t-2} \ \dots \ y_{t-I}]^T$ as a vector denoted by \mathbf{x}_t . Using this notation, the model in (1) is written in a compact matrix form:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad (2)$$

where $\mathbf{y} = [y_t \ y_{t+1} \ \dots]^T$, \mathbf{X} is a matrix where t^{th} row is $[1 \ \mathbf{x}_t^T]$, $\boldsymbol{\beta} = [\beta_0 \ \beta_1 \ \dots]^T$, and $\boldsymbol{\varepsilon} = [\varepsilon_t \ \varepsilon_{t+1} \ \dots]^T$. Vectors \mathbf{y} , $\boldsymbol{\beta}$ and $\boldsymbol{\varepsilon}$ have dimension T . Matrix \mathbf{X} has P columns, which we refer to as the dimension of \mathbf{X} .

Applying standard OLS to (2), the estimate of $\boldsymbol{\beta}$ is given by:

$$\hat{\boldsymbol{\beta}}_{OLS} = \arg \min_{\boldsymbol{\beta}} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2. \quad (3)$$

Under some standard assumptions, the OLS estimator $\hat{\boldsymbol{\beta}}_{OLS}$ is a consistent estimator for the true $\boldsymbol{\beta}$, meaning that the expected difference between the estimator and the true value

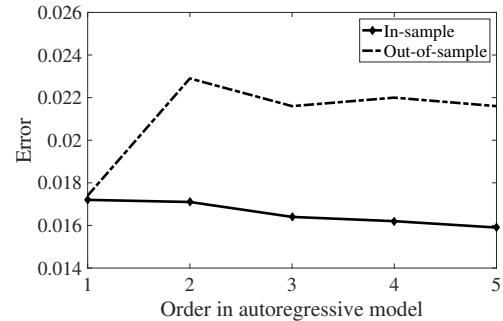


Fig. 1: Error comparison between different autoregressive models.

approaches zero when sample size becomes large [40]. This means that the bias goes to zero as the sample size becomes big. The other asymptotic analysis on the OLS also applies in this case, such as asymptotic distribution for gaussianity of estimators and significance tests; however, in autoregressive models the estimators are typically not unbiased.

To learn and evaluate the estimator $\hat{\boldsymbol{\beta}}_{OLS}$ from OLS, we separate the dataset into a training set and a test set. The estimator $\hat{\boldsymbol{\beta}}_{OLS}$ is learned from the data from the training set and the estimation error is evaluated on the test set. Note that estimators may exhibit extremely good fits on the training set but poor estimation performance on the test set, as according to [41].

III. SPARSITY IN AUTOREGRESSIVE MODELS

Since the objective of OLS estimators is to minimize the sum of squared errors in the training set, OLS achieves optimal *in-sample* performance. This means that adding more regressors into (2) can always decrease the sum of squared error and better fit the data within the training set. However, when we include too many irrelevant regressors, i.e. when we include too many lag orders from the historical data in (2), we are misled by the reduced in-sample bias. We will then ignore the high variance introduced by the estimator which leads to model overfitting.

We use the PG&E dataset as an illustrative example. In this particular dataset, hourly consumption data for single households is recorded. If we use an AR(5) model for a particular household, it will result in an average in-sample squared error of 0.0159, with an average out-of-sample error of 0.0216, whereas AR(1) model has an average in-sample squared error of 0.0172, together with an average out-of-sample error of 0.0174. Thus AR(1) gives better out-of-sample fitting results. If the potential lag orders are up to 10 days, i.e., 240, then an AR(240) model would produce large out-of-sample errors. Overall, we need to select the lag orders carefully to avoid model overfitting. The tradeoff between out-of-sample error and the order of the autoregressive model is shown in Fig. 1.

Determining the correct lag orders is not trivial because it is a combinatorial problem, which is NP-hard. To this purpose, we use LASSO, which is a convex relaxation of such combinatorial problems, to select relevant lag orders. The

intuition for using LASSO is to get a sparse autoregressive models with high orders.

More formally, LASSO is a shrinkage and selection method for linear regression described in (2) [11]. It shrinks parameter estimates towards zero in order to avoid overfitting as well as picks up the most relevant regressors. Operationally, it minimizes the usual sum of squared errors, with a bound on the sum of the absolute values of the coefficients:

$$\hat{\boldsymbol{\beta}}_{\text{LASSO}} = \arg \min \frac{1}{2} \|(\mathbf{y} - \mathbf{x}\boldsymbol{\beta})\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1, \quad (4)$$

where λ is a tuning parameter to control the level of sparsity in the solution. In essence, it controls the number of past consumptions included in the prediction of future consumptions (i.e. sparsity). The bigger λ is, the more sparse the solution $\hat{\boldsymbol{\beta}}_{\text{LASSO}}$ is, which means that less number of data points from the past is used to predict the future. This parameter balances the training and testing performances of the model. When $\lambda = 0$, the solution is the same as in (3) as a traditional autoregressive model. For practical purposes, we are using k -fold cross validation (CV) to determine the value of λ in our simulations, where k is either 5 or 10 [41].

LASSO has gained wide spread popularity in signal processing and statistical learning, see [42], [43], [44]. LASSO has also been applied to forecast electricity price [20], [45], but its application to load forecasting is still a new topic. In [46], LASSO has been applied to forecast short and middle term load. In our paper, we adopt LASSO to predict hour-ahead consumption and provide a comparison between LASSO and some other well-used prediction methods in literature. As discussed later in Section V, LASSO is the most preferable model when considering model simplicity and prediction performances. Despite the fact that RF achieves the best prediction performance with an average relative error less than 20%, it is a highly non linear model and requires to tune many hyper-parameters. This is computationally expensive and not linear in nature. On the other hand, LASSO achieves similar performance as RF and outperforms all other linear models considered in this paper, i.e., SVM with a linear kernel, linear regression with PCA, AR(1) and simple averaging method. It reduces the relative error by 30% compared to these models. Thus LASSO is considered the best in terms of simplicity and prediction accuracy based on our dataset.

Another advantage of using LASSO to recover sparsity is that it has some nice properties as to sign consistency [12], [13]. This means that $\hat{\boldsymbol{\beta}}$ has the same support as $\boldsymbol{\beta}$ and the sign of each element in this support is recovered correctly. Therefore LASSO recovers the exact sparsity of the underlying model. In our simulation, LASSO selects both the most recent lag orders and lag orders with intervals of roughly 24 hours, which performs as a combination of simple averaging and AR(.). Furthermore, LASSO also gives more intuitive results with respect to selected orders. In our simulation, for one electricity user as an example, LASSO selects lag orders as 1, 2, 5, 6, 16, 22, 23, 24, 48, 143, 144, 160, 191, 216, 238, 240. From these orders we can observe a clear behavior pattern of an interval of 24 hours. Some are multiples of 24 and some are not but close to multiples of 24. We thus observe that

not every lag order that LASSO picks is a multiple of 24, otherwise we would directly employ simple averaging rather than LASSO, so LASSO is more adaptive and flexible than simple averaging or AR(.). This implies that user behavior at current hour depends on similar hours happened in previous days. Unlike simple averaging which fix the lag orders at 24, 48, 72, etc., LASSO will automatically select these orders for each individual based on their respective historical data, instead of directly imposing fixed orders.

As can be seen from Section V, LASSO applied to autoregressive model achieves the best prediction performance among all the linear models considered in this paper, with a relative error as small as 22.5 %. It improves the prediction performance by 33.6 % as compared to an AR(1) model.

IV. USER PAIRING BY SIGNIFICANCE TEST

So far we have considered using historical data of an individual user for its own prediction. One way to leverage the fact that we have many users' data is to improve the univariate autoregressive model by including other user's historical data into the model. One popular way to perform this is to conduct Vector Autoregression (VAR), which extends the univariate autoregressive model to joint prediction for a vector of time series data.

To perform a complete VAR, we need to include all potentially relevant users into the autoregressive model, which will reduce the bias but increase the variance for estimators. This causes the same overfitting problem as occurred in univariate autoregressive model, when AR(3) yields a worse prediction on the test set but a better fit on the training set compared to AR(1). One possible way to overcome this problem is to first cluster similar users together and then perform VAR for each cluster. However, consider a scenario where two time series have the exact same values for each time slot. Then these two time series are clustered together since they are identical. In this case, clustering them together and performance multivariate autoregressive model does not help to enhance prediction because knowing the history of one time series would not help predict the future values of another time series. This problem distinguishes similarity from prediction performance, which is the focus of this section.

In this paper we focus on selecting the most relevant user to enhance prediction performance after doing univariate LASSO selection. To this end, we adopt LASSO significance test to select this most relevant user. In LASSO significance test, the inclusion of a particular user is based on how well this particular user's history data explains the fitted residual after performing LASSO to one user's univariate autoregressive model. This is a hypothesis test of an exponential random variable [47]. We will discuss more details of a LASSO solution to a regression model and the implementation of the LASSO significance test in the rest of this section. Overall from the simulation results presented in section V, using significance test on top of univariate LASSO-type regression model improves the relative prediction error from 22.5% to 20.9%, which is almost as good as the prediction results from RF.

A. Linear Regression Model for the Fitted Residual and LASSO Path

In LASSO significance test, we want to test if the fitted residual from LASSO solution in (4) is indeed noise or if it can be better explained by other user's historical data. Intuitively, We need to test inclusion of each user's historical data. We therefore generate LASSO autoregressive coefficients for univariate time series in (1) and compute the residual $\{e_t\}$ (true value minus fitted value) for each individual user. Then we use e_t as the response variable in a new linear regression model. To fit this new linear regression model, we include all users' historical data at lag order one except for this user as the regressors. In this way we will have a high dimensional regressor matrix. Its dimension is the number of users we want to potentially include, which is the total number of users minus one. Since we only want to include the most significant user as the regressor, LASSO is therefore performed and furthermore tested on this new linear regression problem. In the following, we first illustrate more details on formulating the LASSO regression model to the fitted residual from univariate model. Then we provide more explanations on the LASSO path with a varying λ . In the next subsection we introduce the covariance statistic and its asymptotic distribution for hypothesis tests along the LASSO path.

Mathematically, the new linear regression problem for the residual is formulated as follows:

$$\mathbf{e} = \boldsymbol{\xi} \boldsymbol{\alpha} + \boldsymbol{\psi}, \quad (5)$$

where $\mathbf{e} = [e_{t,s} \ e_{t+1,s} \ \dots]^T = [y_{t,s} - \mathbf{X}_{t,s}^T \hat{\boldsymbol{\beta}}_{LASSO} \ y_{t+1,s} - \mathbf{X}_{t+1,s}^T \hat{\boldsymbol{\beta}}_{LASSO} \ \dots]^T$ is the residual for the current user s , $\mathbf{X}_{t,s} = [y_{t-1,s} \ y_{t-2,s} \ \dots \ y_{t-I,s}]^T$ and $y_{t,s}$ denotes the consumption data for user s at time t . The parameter $\boldsymbol{\alpha}$ is a vector of weights that denotes how helpful each other user's last demand data is towards the prediction of future demand of user s . Vector \mathbf{e} has a length of N . Moreover, matrix $\boldsymbol{\xi}$ is the regressor matrix made by lag order one historical data from all the other users except for the current user s . The t^{th} row of $\boldsymbol{\xi}$ is $[y_{t-1,1} \ y_{t-1,2} \ \dots \ y_{t-1,s-1} \ y_{t-1,s+1} \ \dots]$. So $\boldsymbol{\xi}$ has a column length of P , which is the number of users to be included, and a row length of N , which is the same as that of \mathbf{e} . One user is represented by one column of $\boldsymbol{\xi}$ and is regarded as one regressor variable. In addition, $\boldsymbol{\psi}$ is the white noise vector with variance σ^2 . The parameter $\boldsymbol{\alpha} = [\alpha_0 \ \alpha_1 \ \dots \ \alpha_P]^T$ is the decision variable for the regression problem.

We again apply LASSO for (5), to avoid overfitting by including too many irrelevant users. Building on the discussion in Section III, we define the LASSO path as the revolution of the estimator $\hat{\boldsymbol{\alpha}}$ in terms of a sequence of λ_k 's. The LASSO path $\hat{\boldsymbol{\alpha}}(\lambda_k)$ is given by:

$$\hat{\boldsymbol{\alpha}}(\lambda_k) = \arg \min_{\boldsymbol{\alpha}} \frac{1}{2} \|\mathbf{e} - \boldsymbol{\xi} \boldsymbol{\alpha}\|_2^2 + \lambda_k \|\boldsymbol{\alpha}\|_1, \quad (6)$$

where λ_k is called the knot along the LASSO path.

For different values of λ_k , we obtain different solutions and sparsity at different levels. The active set at one particular value of λ_k is the set of all non zero coefficients estimated at that value, i.e., $\mathcal{A} = \{\hat{\alpha}_p \neq 0, \lambda = \lambda_k, p = 0, 1, \dots, P\}$. The

path $\hat{\boldsymbol{\alpha}}(\lambda_k)$ is continuous and piecewise linear with knots at these values $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$ [48]. With the formulation in (5), the goal is to test if an inclusion of one user's historical data is helpful for prediction. Mathematically speaking, we want to test if the variables that sequentially enter the active set are statistically significant.

B. Covariance Test

Significance test applied to the LASSO path is fundamentally different from the standard hypothesis tests in linear regression. In standard testings with linear regression model, we fix a regressor variable and test the significance of its inclusion into the model. This can be done by t -test. However, the scenario is not the same in the case where we want to test the significance of variables along the LASSO path since the variables entering the active set are not known in advance. This means that we do not know which regressors to fix in order to conduct t -test. This problem is addressed by the authors in [47]. Instead of standard t -test, they propose a significance test for regressors progressively selected by the LASSO path, using the so called covariance statistic. They have shown that the covariance statistic asymptotically follows exponential distribution and can be used to test the significance of the entry of variables into the active set.

Using the covariance test defined in [47], we test whether it is significant to include one variable (one other user) into the active set against including no variables into the active set. Specifically, we want to analyze whether the residual, i.e., \mathbf{e} is white noise, or it can be explained significantly by the lag one order historical data from other users.

The null hypothesis is that the true active set is an empty set, i.e., no other users can help better predict the current user:

$$H_0 : \alpha_1 = \alpha_2 = \dots = \alpha_P = 0. \quad (7)$$

Another way to interpret this is:

$$H_0 : \mathbf{e} \stackrel{H_0}{\sim} N(0, \sigma^2 I). \quad (8)$$

We focus on testing the significance of the first variable entering the active set and the two main reason for doing so are as follows. First, for the following variables entering the active set, the exponential distribution turns out to yield a slightly higher value than the true value of the test statistic, so the decision tends to be conservative [47]. Second, If one would like to include more significant regressors selected by LASSO, the Tailstop criterion [49] for ordered selection can be used under the assumption that sequential p -values are independent. However, we restrain from looking at p -values at greater indices because of the non-orthogonality of $\boldsymbol{\xi}$ in our case (which leads to misleading interpretation of p -values), as pointed out in [49] and [50]. To illustrate this, we use Tailstop criteria in our simulation and it turns out that the first eighteen regressors are selected to limit the type I error of the global null in (7) under 0.05. It is obviously an overfitting result and Tailstop criterion cannot be trusted due to the correlation between each row in the regressor matrix $\boldsymbol{\xi}$. From literature, interpreting p -values from sequential hypothesis and order

selection under generic conditions is still an open question and is thus beyond the discussion of this paper.

To test whether the first variable entering the active set is significant, we first set $\lambda = \infty$ and gradually reduce λ until one regressor variable has a non-zero α . Denote the value of this λ as λ_1 . Also denote λ_2 as the value of λ when the second regressor variable enters the active set. For simplicity of representation, we scale the columns of the ξ so that each column has unit norm.

Following this scaling strategy, the authors in [47] define the covariance statistic for testing the entry of the first variable and it is written as:

$$T_1 = \frac{\lambda_1(\lambda_1 - \lambda_2)}{\sigma^2}. \quad (9)$$

From Theorem 2 in [47], the authors have proved that $\Pr\{T_1 > t\} \rightarrow e^{-t}$ as $P \rightarrow \infty$ for all $t \geq 0$. Thus we have that the asymptotic distribution of T_1 obeys exponential distribution with parameter 1: $T_1 \rightarrow \text{Exp}(1)$.

In most cases, the value of σ^2 is not known a priori. We can estimate σ^2 by the residual sum of squared error via:

$$\hat{\sigma}^2 = \frac{\|(\mathbf{e} - \xi \hat{\alpha}_{OLS})\|_2^2}{N - P}, \quad (10)$$

where $\hat{\alpha}_{OLS}$ is obtained through:

$$\hat{\alpha}_{OLS} = \arg \min_{\alpha} \frac{1}{2} \|(\mathbf{e} - \xi \alpha)\|_2^2. \quad (11)$$

Plugging (11) into (9), we have a new statistic F_1 , which is asymptotically following F distribution (ratio of two independent χ^2 distribution) under the null [47]:

$$F_1 = T_1 \frac{\hat{\sigma}^2}{\hat{\sigma}^2} = \frac{\lambda_1(\lambda_1 - \lambda_2)}{\hat{\sigma}^2} \rightarrow F_{2, N-P}. \quad (12)$$

In conclusion, to test whether another user's historical data can explain the residual obtained by the univariate LASSO introduced in Section III, we compute the value of λ_1 and λ_2 , along with the full linear regression in (11) which gives us $\hat{\sigma}^2$. Plugging these into (12), we include the regressor variable (which represent one particular user) entering the active set at λ_1 if F_1 is greater than some threshold and reject the null hypothesis in (8). Simulation results show that performing this significance test on top of univariate LASSO type regression model has improved the prediction by reducing relative test error by 38.3%, compared to AR(1) model.

V. SIMULATION RESULTS

We use the data from PG&E. It contains hourly smart meter readings for residential users during a period of one year from 2010 to 2011. Temperature data is retrieved from an online database [51] for the same period. Some preliminary observations suggest that temperature is not a significant regressor to predict consumption based on our data set, thus it is excluded from the regressors for the following simulations. In addition, the consumption data has been filtered to exclude daily periodic trends and is separated into weekday data and weekend data. The simulation results are obtained on weekday data.

A. Prediction Evaluation Criteria

Before proceeding with the prediction methods, we first introduce the evaluation criteria for prediction used in this paper. A naive way to evaluate prediction methods is to compare the Mean Squared Error (MSE) or Mean Absolute Error (MAE) within the testing set:

$$\text{MSE} = \frac{1}{n} \sum_{t=1}^n (y_t - \hat{y}_t)^2 \quad (13)$$

$$\text{MAE} = \frac{1}{n} \sum_{t=1}^n |y_t - \hat{y}_t| \quad (14)$$

where y_t is the actual value, \hat{y}_t is the predicted value, and n is the number of fitted values.

However, MSE and MAE are not scale invariant. This will be misleading if we want to compare the performance for two datasets with significantly different scales and means. One way to solve this problem is to normalize the error or to compare the relative error, i.e., the prediction error with respect to the data scale. Here we use the MAPE mentioned earlier in Section I, to capture the relative error:

$$\text{MAPE} = \frac{1}{n} \sum_{t=1}^n \left| \frac{y_t - \hat{y}_t}{y_t} \right|. \quad (15)$$

Normalized Root-Mean-Square Deviation (NRMSD) is also a good metric to compare prediction performances of data with different scales :

$$\text{NRMSD} = \frac{\sqrt{\frac{1}{n} \sum_{t=1}^n (y_t - \hat{y}_t)^2}}{\bar{y}}, \quad (16)$$

where \bar{y} is the mean of y in the training set.

In case of outliers, we adopt mean curtailing of 0.01 tail and head, or simply use the median to replace the mean value.

In addition, to evaluate the model between good fit and complexity (in terms of the number of parameters), Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) can be used [52].

B. Comparison with Different Prediction Methods

We compare our proposed method with other popular methods: 1): Simple averaging, which averages the consumption at the same hour of the day for ten previous days; 2): LW, which uses the consumption of the same hour of the day during last week; 3): AR(1); 4): ES with additive error; 5): linear regression with PCA; 6): SVM with and without PCA; 7): RF and 8): NN with a single hidden layer. Input features are ten previous days' consumption data which has a dimension of 240, except that for NN the input features are set to be previous day's consumption data, which has a dimension of 24. While the first and the second method are very intuitive methods for prediction, AR(1) and ES are popular prediction methods in analyzing time series data. The remaining methods are popular machine learning algorithms and artificial intelligence techniques in prediction. Particularly RF and NN are the only non linear methods considered in this paper. They serve to provide reference results as compared to other linear methods. Since the dataset that we use only contains time series load

consumption data (plus coarse zip code data), we restrain from using more complexed models which tend to overfit. In all, our simulations using real data show that our proposed method performs similar or better to these sophisticated algorithm while remaining a simple linear structure.

We make the following clarifications for the machine learning techniques in 5)-8). First, for the extracted features obtained from PCA, we threshold the number of components with a tolerance ζ on how much the chosen components explain the covariance matrix of data ξ . Then we include those components as new features to the linear regression model. If we set up a tolerance $\zeta = 0.75$ and omit components if their standard deviations are smaller than ζ times the standard deviation of the first component, we obtain an average two principle components for each user based on our dataset. Second, a linear kernel is trained as we are comparing SVM with several linear regression models. Third, for RF, we grow 500 trees in total. At each split of the tree node we re-sample a third of the original features. Note that RF is a nonlinear model in terms of input features whereas the other models are linear in input features. For a more in-depth introduction to RF, please refer to [53]. Last, due to convergence and computational issues, we consider one hidden layer with five neurons in the NN model.

We visualize the comparison in terms of MAPE in Fig. 2, with 50 users in the training set. In the simulations, the forecast horizon is one hour and the shown errors are the averaged medians over the users. Detailed comparison results based on MAPE, MAE, MSE and NRMSE are presented in Table I, Table II, Table III and Table IV. In Table I to Table IV, the subscript W denotes the window size. Result from linear regression after PCA is outside the range of the results from the other prediction methods shown in Fig. 2 and is not presented. In addition, there is a constraint for the minimum number of samples in order to have a well behaved LASSO estimator. This lower bound is of order $O(s \log p)$ [13], where s is the cardinality of the true support and p is the total number of regressors. However, since we do not know the level of sparsity s in advance, the number of samples should be at least $O(p \log p)$, which is around 600. Nevertheless, since the columns in the regressor matrix are correlated, we resort to a training size larger than the theoretical bound. In the mean time, we cannot include too many samples into the training set to avoid losing stationarity. Here unit root test is applied to test stationarity. To these two ends, we experiment the training sizes of 720, 960 and 1200 samples and compare the respective performances by the proposed prediction methods.

From Fig. 2, we see that RF yields the smallest prediction error in terms of MAPE. This is not unexpected since RF is a complex method that is highly nonlinear. Surprisingly, our proposed LASSO method matches the performance of RF when window sizes are increased. With a sliding window with length of 720 hours, RF has an average MAPE of 0.211 and LASSO has an average MAPE of 0.231. When we increase the size of the sliding window to include 1200 hours' consumption data, RF has an average MAPE of 0.203 and LASSO has an average MAPE as 0.206. At this stage, operators may prefer to use LASSO since it is as accurate as RF but has

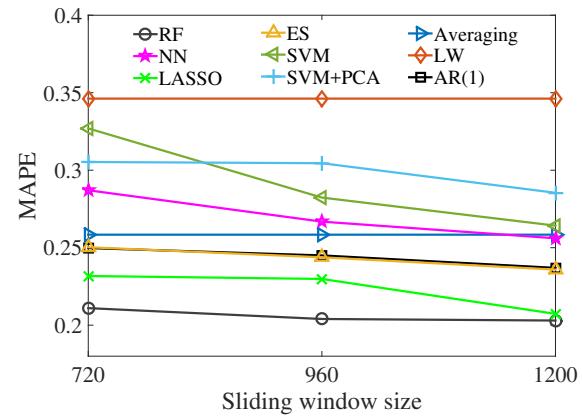


Fig. 2: MAPE for different prediction methods with various window size of the training set. RF yields the best prediction performance among all aforementioned methods and LASSO yields the best performances among linear models.

TABLE I: Average MAPE cross the 50 users from the dataset for the prediction methods. Subscript W denotes the window size.

Method	MAPE _{W=720}	MAPE _{W=960}	MAPE _{W=1200}
Averaging	0.258	0.258	0.258
LW	0.346	0.346	0.346
AR(1)	0.249	0.245	0.237
ES	0.250	0.244	0.236
SVM	0.327	0.282	0.264
SVM after PCA	0.305	0.304	0.285
RF	0.211	0.204	0.203
NN	0.287	0.267	0.256
LASSO	0.231	0.229	0.206

TABLE II: Average MAE cross the 50 users from the dataset for the prediction methods. Subscript W denotes the window size.

Method	MAE _{W=720}	MAE _{W=960}	MAE _{W=1200}
Averaging	0.0373	0.0373	0.0373
LW	0.0540	0.0540	0.0540
AR(1)	0.0322	0.0318	0.0311
ES	0.0340	0.0334	0.0322
SVM	0.0393	0.0338	0.0322
SVM after PCA	0.0363	0.0357	0.0368
RF	0.0273	0.0265	0.0260
NN	0.0360	0.0336	0.0311
LASSO	0.0291	0.0290	0.0260

much lower model complexity. For example, using LASSO operators can understand the most significant hours in the past that influence the prediction. In addition, the coefficients in LASSO can be understood as sensitivities of future prediction on past consumption, due to the linear structure of LASSO. In contrast, the coefficients in RF do not have such "causal" structure. In addition, LASSO can be computed much faster than RF.

Besides these two prediction methods, SVM with either

TABLE III: Average MSE cross the 50 users from the dataset for the prediction methods. Subscript W denotes the window size.

Method	$MSE_{W=720}$	$MSE_{W=960}$	$MSE_{W=1200}$
Averaging	0.00143	0.00143	0.00143
LW	0.00231	0.00231	0.00231
AR(1)	0.00103	0.00101	0.00096
ES	0.00113	0.00105	0.00094
SVM	0.00155	0.00115	0.00104
SVM after PCA	0.00132	0.00128	0.00125
RF	0.00074	0.00070	0.00067
NN	0.00123	0.00108	0.00097
LASSO	0.00085	0.00084	0.00067

TABLE IV: Average NRMSE cross the 50 users from the dataset for the prediction methods. Subscript W denotes the window size.

Method	$NRMSE_{W=720}$	$NRMSE_{W=960}$	$NRMSE_{W=1200}$
Averaging	0.197	0.197	0.197
LW	0.241	0.241	0.241
AR(1)	0.166	0.156	0.135
ES	0.168	0.153	0.125
SVM	0.221	0.182	0.153
SVM after PCA	0.200	0.194	0.173
RF	0.157	0.143	0.126
NN	0.178	0.167	0.165
LASSO	0.172	0.163	0.126

the original features or the reduced features are average in their prediction performances. NN does not outperform other methods. We believe this is due to the lack of data. Since we only have time series consumption data without any other side information, only a shallow NN (a single hidden layer) is trained to avoid overfitting and other computational issues.

The nonparametric methods, LW and simple averaging method, return higher MAPEs compared to machine learning techniques. However, based on our dataset, simple averaging using ten previous day's data still outperforms LW (using data from last week only) because the noises are leveraged by adding more consumption data into account. Since these nonparametric methods do not depend on the size of the sliding window, the MAPEs are the same across different sliding window sizes. What is more, these nonparametric methods are powerful when the training set is small (a smaller sliding window size). However, when more data are available to be included into the training set, the performances of these methods do not get improved.

In addition, the time series analysis methods, namely AR(1) and ES, yield moderate MAPE. However, it is still not as good as either RF or LASSO. In conclusion, LASSO achieves as good as the performance from RF, and other methods return inferior prediction performances. Details are shown in Table I. Table II, Table III to Table IV summarize more results of the comparison based on different metrics. In all, LASSO is shown to be consistently the among the best methods and can provide performances that are as good as advanced machine

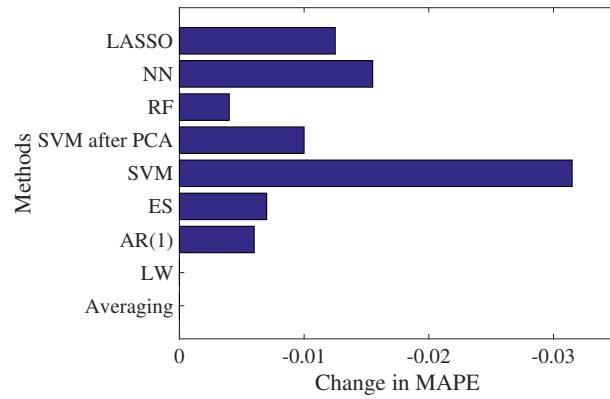


Fig. 3: Average change of MAPE with an increase of window size.

learning techniques.

Besides the above observations on predictive performance of various models, we can also see a difference in model sensitivity with respect to an increase in window size (a larger training set). The results are shown in Fig 3, in terms of MAPE.

As can be seen from Fig 3, LW and averaging have no sensitivity to window size. This is because they are both an averaged sum of a fixed history and are not parametric. AR(1) and ES have similar sensitivities since they are both autoregressive models. SVM originally has a high sensitivity to window size because we include many features into the model and fix a linear kernel. However, this may not capture correctly the underlying consumption behavior. However, with dimension reduction from PCA, the sensitivity from SVM is improved. Similarly, NN and LASSO do not have a negligible sensitivity towards window size because of the large number of features in the model but since they are more powerful models than SVM, the sensitivity is not as high as the latter. In addition, RF falls into the scheme of ensemble learning algorithm and its variance should be smaller than the other parametric models. Therefore its sensitivity towards window size is small.

We then compare the averaging computational time for the aforementioned methods. The simulations are tested on a MacBook with 2.7 GHz Intel Core i5 processor and 8GB 1867 MHz DDR3 memory. The empirical computational time for each prediction method is shown in Table V.

As can be seen from Table V, Simple averaging, LW, AR(1) and ES are most time efficient but at the sacrifice of prediction accuracy. Linear regression with PCA is moderately efficient but provides the worst prediction accuracy. RF yields the worst computational time because it needs to be trained on the repeated bootstrapping samples. NN is also not preferable because for a single layer model it already takes up to more than 14 seconds to train over a single sliding window. LASSO beats both SVM with either the original regressor matrix or the principle components in terms of time efficiency. Combining the observations from Table V with Fig. 2, we can see that for univariate time series analysis, LASSO best trades off between

prediction accuracy and time efficiency.

Based on the observations from Fig. 2 and Table V, the complexity grows as the prediction model varies from simple averaging method up to RF and consequently the computational time grows accordingly. It is even more time consuming when we try to tune the hyper-parameters in RF, for example, the number of features to bootstrap at each split of the tree, total number of trees, etc... What is more, although RF provides the most accuracy, it is the hardest to be understood in terms of the estimated parameters from the model, i.e., the decision boundary at each split of the tree node for each tree. Since throughout this paper the focus is on one hour ahead prediction, this complexity ruins the performance of the model and is not preferred. In contrast, LASSO yields a competing prediction performance with RF and has only one hyper-parameter (λ) to tune. The optimal value of this hyper-parameter can be determined from cross-validation and is easy to implement. Thus it is the most preferable method for univariate analysis based on our dataset.

Table VI further consolidates this preference in terms of complexity of input data. Feature dimension is defined as the number of input variables in the model. Simple averaging method has a dimension of 10 because it is a average of consumption data during the same hour for the past ten days. LW only needs one data point during last week to make prediction. AR(1) and ES also have a much smaller input space but their prediction performances are not as good as either LASSO or RF. For methods associated with dimension reduction by PCA, the feature dimension is the number of the reduced features for the model, thus around 20 in our case. We restrict the feature dimension as 24 for NN since a higher feature dimension will require more iterations in the algorithm to reach convergence. Other than that, for SVM, RF and LASSO, the input data has a dimension of 240 which stands for the total number of hours in the past ten days. With the same input complexity, LASSO yields the most efficient computational time from Table V.

Based on the discussions from Fig. 2, Table V and Table VI, LASSO is the best choice in terms of model efficiency and model simplicity.

TABLE V: Average computational time per user per sliding training window for each prediction method.

Prediction method	Time (in seconds)
Averaging	less than 0.0001
LW	less than 0.0001
AR(1)	0.005
ES	0. 093
PCA + linear regression	0.175
SVM	9.347
SVM + PCA	9.522
RF	21.562
NN	14.416
LASSO	0.220

In addition, the average computational time for computing the optimal pairs using the covariance test is 3.51 seconds on average per each user. We then conclude that it is time efficient

TABLE VI: Feature dimension for different prediction methods.

Prediction method	Feature dimension
Averaging	10
LW	1
AR(1)	1
ES	at most 3
PCA + linear regression	around 20
SVM	240
SVM + PCA	around 20
RF	240 (split the tree 500 times.)
NN	24
LASSO	240 (active regressors around 10-20.)

to compute the optimal pairs of users. What is more, since the optimal pair assignment is fixed once we compute it, this computational time can be amortized and is nearly negligible.

C. Results for Univariate Time Series Analysis

Based on the comparison result shown in Section V-B, in this subsection we perform more detailed simulations on the following prediction methods from the methods considered so far:

- Simple averaging method, which is a representative method from nonparametric methods;
- AR(1), which is a representative method from time series methods;
- RF, which is a representative method from machine learning techniques;
- LASSO applied in autoregressive model, which is the proposed method in this paper.
- SVM with and without PCA, which we use as an illustration of the failure of PCA based on our dataset.

The simulation results associated with the covariance test are going to be introduced later in Section V-D. The results of the prediction methods besides covariance test are shown in Fig. 4 and Table VII, for a random selection of 150 users from the dataset. We use MAPE as the criterion to compare prediction performances because it is a normalized metric. Again, the MAPE shown is the median over all the sliding windows and the forecast horizon is one. As can be seen in Fig. 4, LASSO type regression reduces the variance of MAPE for these 150 users in the dataset. Its mean is also reduced. A comparison of the aforementioned methods is summarized in Table VII.

From Fig. 4, we see that MAPE varies on a user to user basis. For example, from Fig. 2(a), using AR(1) model to predict consumption will yield a MAPE greater than 50% for some of users. What is more, there is one extreme case for simple averaging method shown in Fig. 2(c) that MAPE reaches 200%. This extreme value of MAPE suggests a rather unfavorable prediction result, even worse than simply predicting zero (which has a MAPE as 100%). However, for both methods, a MAPE close to zero can still be achieved for other users. Besides these methods, SVM achieves a moderate deviation of MAPE but it is still higher than by LASSO and RF.

Comparing Fig. 2(a), Fig. 2(c) with Fig. 2(b), Fig. 2(d), Fig. (e) and Fig. (f), we see that LASSO and RF have MAPEs that are less spread between 0% and 100%. In addition, LASSO and RF have a smaller mean of MAPE than that from the other prediction methods. More detailed statistics are shown in Table VII. What is more, comparing Fig. 2(e) and Fig. 2(f), PCA does not improve the performance of prediction models such as SVM, as we have discussed in Section V. B. Here, prediction by SVM using the original features (historical consumption data) as inputs achieves a smaller MAPE than that after dimension reduction by PCA. SVM after PCA has a MAPE of 0.291 against a MAPE of 0.264 by SVM alone.

In addition, besides the improved prediction performance compared to AR(1) and simple averaging method, LASSO also provides some straightforward understandings of the model. Particularly, we include as much as 240 lag orders for LASSO, which includes all the historical data for the previous ten days. Taking user No.1 as an example, LASSO selects 16 non zero lag orders, according to 10-fold CV with a sequence of decreasing $\{\lambda_k\}$. The lag orders that LASSO picks are 1, 2, 5, 6, 16, 22, 23, 24, 48, 143, 144, 160, 191, 216, 238, 240. This pattern reflects that LASSO not only selects the most recent lag orders (which is similar to AR(1)), but the lag orders roughly at interval of one day, i.e., 24 hours as well (which is similar to simple averaging). The coefficients for the lag orders also have different scales. In the example of user No.1, the most recent lag orders given the largest coefficient ($\beta_1 = 0.259$). The second largest coefficient is given to lag order 24 ($\beta_{24} = 0.187$). The rest of the coefficients are scaled between 0.01 to 0.06.

Note that LASSO will return selected lag orders adaptively for different users. For user No.10, the lag orders that LASSO picks up for this user are 1, 23, 24, 120, 138, 144, 168, 239, 240. In this case, the three most recent lag orders are given the largest coefficients, where $\beta_1 = 0.579$, $\beta_{23} = 0.123$ and $\beta_{24} = 0.103$. The rest of the coefficients scale at around 0.02. So for this particular user, LASSO picks up less historical data at the same hour than user No. 1, at roughly one day ago, five days ago, six days ago, seven days ago and ten days ago. This is different from simple averaging, since LASSO does not pick up the lag order 48 or 72, which are more recent lags than 120. Still, however different lag orders that LASSO picks for each user, we can observe a clear pattern of multiples of 24 hours, indicating user's consumption behavior.

In all, the proposed model based on LASSO is useful in practice because of its simplicity and linear nature. Given the set of coefficients, utilities and planners can easily identify and attribute the impact of different features on load consumption. For example, suppose for a given user, our method concludes that its load at time t is mostly determined by the past load at times $t-1$, $t-3$ and $t-24$. This shows that this user has some short term behavior in the one hour range, medium term behavior in the 3 hour range, and a daily cycle partner that repeats every 24 hours. This information can be interpreted by utilities to better understand this user's consumption patterns and potentially identify appropriate demand programs.

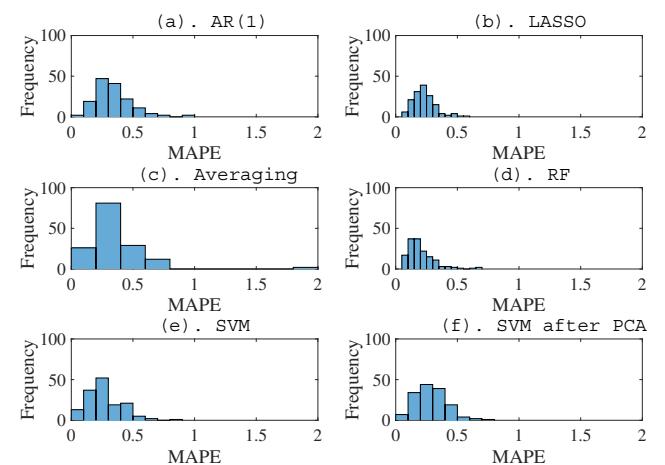


Fig. 4: Histogram comparison for univariate time series. Mean of MAPE for AR(1) is 0.339, for LASSO 0.225, for simple averaging method 0.359, for RF 0.195 and for SVM with and without PCA is 0.264 and 0.291 respectively. Standard deviation of MAPE for AR(1) is 0.151, for LASSO 0.094, for simple averaging method 0.237, for RF 0.109, for SVM 0.142 and for SVM after PCA is 0.129.

D. Results for Pairing the Users

As described in Section IV, we include historical consumption data from other users to fit the current user's data. We compare the results of the two methods discussed as follows:

- 1) Reference method: LASSO selection in autoregressive model using one user's own data.
- 2) Pairing by covariance test: the optimal pairing for a user is selected by the best user fitted to the residuals obtained from this user's consumption data after LASSO in univariate autoregressive model. This best user corresponds to the first regressor in ξ entering the active set when its associated p -value is small. The lag orders for this user are selected by LASSO in the reference method, but the lag order for its paired user is fixed to one.

The results for these two methods are in Fig. 5 and Table VII, with a forecast horizon as one. The results are based on the same 150 users chosen from the dataset. It can be seen from Fig. 5 that with the covariance test that includes another useful user into the model, the prediction performance is enhanced. The mean of the MAPE on the test set is improved by 7% and the variance by 6%. Note that the selection of another user into the model is not mutual, for example, for user No.14, user No.137 is selected by covariance test whereas for user No.137, user No.50 is selected. This is essentially different from bivariate autoregressive model where pairs of users are included into the model simultaneously. Here, the pairing of the users is adaptive according to the covariance statistic defined in (9).

As can be seen from Table VII, the proposed methods greatly improve the prediction performance by reducing the relative prediction error, thus will provide more reliable results on consumption prediction. The proposed prediction methods

TABLE VII: Results for the prediction methods, where CovT stands for covariance test and Ave stands for simple averaging method.

	Ave	AR(1)	LASSO	RF	SVM	CovT
mean (MAPE)	0.359	0.339	0.225	0.195	0.264	0.209
sd (MAPE)	0.237	0.151	0.094	0.109	0.142	0.088

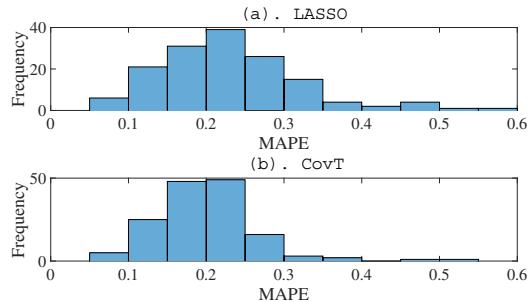


Fig. 5: Histogram comparison for multivariate time series, where covT stands for covariance test. Mean MAPE for LASSO is 0.225 and for performing covariance test after LASSO is 0.209. Standard deviation of MAPE has also reduced by 7% from LASSO to performing covariance test after LASSO.

in this paper can help compare the user behavior with and without demand response incentives. In the future, we will use the results in this paper to investigate the demand response effect of users by comparing their actual consumption data and predicted consumption data.

VI. CONCLUSIONS AND FUTURE WORKS

In this paper, we propose a sparse autoregressive model to predict the electricity consumption of individual users. First, we adopt LASSO to recover the sparsity in linear regression models. LASSO selects the most recent lag orders and important lag orders with close multiples of 24 hours, which reveals user consumption patterns. Second, we improve prediction accuracy of a particular user by leveraging other user's historical data, where we use the covariance statistic to test if inclusion of another user is significant to explain the fitted residual from univariate LASSO-type regression. Third, we give a comprehensive analysis on the proposed methods against several available predictive models in literature. We observe that LASSO best trades off between model complexity and prediction performance. It has the least predictive error among the linear models and yields a competitive performance as compared to ensemble methods such as random forests. Further simulation results show that LASSO with covariance test outperforms both simple averaging method and auto regressive method with order 1 for individual consumption prediction, improving the prediction performance by 38.4% in terms of MAPE.

REFERENCES

- H. Afares and M. Nazeeruddin, "Electric load forecasting: literature survey and classification of methods," *International Journal of Systems Science*, vol. 33, no. 1, pp. 23–34, 2012.
- J. W. Taylor and P. E. McSharry, "Short-term load forecasting methods: An evaluation based on european data," *Power Systems, IEEE Transactions on*, vol. 22, no. 4, pp. 2213–2219, 2007.
- R. Weron, *Modeling and forecasting electricity loads and prices: a statistical approach*. John Wiley & Sons, 2007, vol. 403.
- R. Sevlian and R. Rajagopal, "Short term electricity load forecasting on varying levels of aggregation," *arXiv preprint arXiv:1404.0058*, 2014.
- M. Baran and T. McDermott, "Distribution system state estimation using ami data," in *2009 Power Systems Conference and Exposition*. IEEE, 2009, pp. 1–3.
- C. Chen, J. Wang, Y. Heo, and S. Kishore, "MPC-based appliance shedding for residential building energy management controller," *Smart grid, IEEE Transactions on*, vol. 4, no. 3, pp. 1401–1410, 2013.
- M. F. S. Aman, C. Charalampou, U. Noor, Y. Simmhan, and V. Prasanna, "Empirical comparison of prediction methods for electricity consumption forecasting," University of Southern California, Tech. Rep., 2014.
- J. MacDonald, P. Cappers, D. Callaway, and S. Kiliccote, "Demand response providing ancillary services," Grid-Interop, Tech. Rep., 2012.
- P. Li, X. Guan, J. Wu, and D. Wang, "An integrated energy exchange scheduling and pricing strategy for multi-microgrid system," in *In TENCON 2013-2013 IEEE Region 10 Conference (31194)*. IEEE, 2013, pp. 1–5.
- D. Wang, X. Guan, J. Wu, P. Li, P. Zan, and H. Xu, "Integrated energy exchange scheduling for multimicorgrid system with electric vehicles," *Smart grid, IEEE Transactions on*, vol. 7, no. 4, pp. 1762 – 1774, 2016.
- R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.
- P. Zhao and B. Yu, "On model selection consistency of lasso," *The Journal of Machine Learning Research*, vol. 7, pp. 2541–2563, 2006.
- M. J. Wainwright, "Sharp thresholds for high-dimensional and noisy sparsity recovery using-constrained quadratic programming (lasso)," *Information Theory, IEEE Transactions on*, vol. 55, no. 5, pp. 2183–2202, 2009.
- H. Zou, "The adaptive lasso and its oracle properties," *Journal of the American statistical association*, vol. 101, no. 476, pp. 1418–1829, 2006.
- L. Meier, S. V. D. Geer, and P. Bhlmann, "The group lasso for logistic regression," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 70, no. 1, pp. 53–71, 2008.
- H. Xu, D. Eis, and P. Ramadge, "The generalized lasso is reducible to a subspace constrained lasso," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2013, pp. 3268–3272.
- F. Bunea, J. Lederer, and S. Y., "The group square-root lasso: Theoretical properties and fast algorithms," *Information Theory, IEEE Transactions on*, vol. 60, no. 2, pp. 1313–1325, 2014.
- L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- T. K. Ho, "Random decision forests," in *Document Analysis and Recognition, 1995 Proceedings of the Third International Conference on*, vol. 1. IEEE, 1995, pp. 278–282.
- F. Ziel, R. Steinert, and S. Husmann, "Efficient modeling and forecasting of electricity spot prices," *Energy Economics*, vol. 47, pp. 98–111, 2015.
- T. Hong, "Energy forecasting: Past, present, and future," *Foresight: The international Journal of Applied Forecasting*, vol. 32, pp. 43–48, 2014.
- NewYorkISO, "Caiso demand response resource user guide," NewYorkISO, Tech. Rep., 2007.
- C. Katie, M. A. Piette, C. A. Goldman, and S. Kiliccote, "Statistical analysis of baseline load models for non-residential loads," *Energy and Buildings*, vol. 41, no. 4, pp. 374–381, 2009.
- S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings, "Putting the smarts into the smart grid: a grand challenge for artificial intelligence," *Communications of the ACM*, vol. 55, no. 4, pp. 86–97, 2012.
- S. Aman, Y. Simmhan, and V. K. Prasanna, "Improving energy use forecast for campus microgrids using indirect indicators," in *Data Mining Workshops (ICDMW), 2011 IEEE 11th International Conference on*. IEEE, 2011, pp. 389–397.
- B. J. Chen, M. W. Chang, and C. J. Lin, "Load forecasting using support vector machines: a study on eunite competition 2001," *Power Systems, IEEE Transactions on*, vol. 19, no. 4, pp. 1821–1830, 2004.
- Y. Weng and R. Rajagopal, "Probabilistic baseline estimation via gaussian process," in *2015 IEEE Power and Energy Society General Meeting*. IEEE, 2015, pp. 1–5.
- H. Quan, D. Srinivasan, and A. Khosravi, "Short-term load and wind power forecasting using neural network-based prediction intervals," *Neural Networks and Learning Systems, IEEE Transactions on*, vol. 25, no. 2, pp. 303–315, 2014.

- [29] A. Sarwat, M. Amini, A. Domijan, A. Damjanovic, and F. Kaleem, "Weather-based interruption prediction in the smart grid utilizing chronological data," *Journal of Modern Power Systems and Clean Energy*, vol. 4, no. 2, pp. 308–315, 2016.
- [30] J. B. Fiot and F. Dinuzzo, "Electricity demand forecasting by multi-task learning," *Smart Grid, IEEE Transactions on*, vol. PP, no. 99, pp. 1–8, 2016.
- [31] O. Abedinia, N. Amjadi, and H. Zareipour, "A new feature selection technique for load and price forecast of electrical power systems," *Power Systems, IEEE Transactions on*, vol. PP, no. 99, p. 1, 2016.
- [32] N. Amjadi, "Short-term hourly load forecasting using time-series modeling with peak load estimation capability," *Power Systems, IEEE Transactions on*, vol. 16, no. 3, pp. 498–505, 2001.
- [33] S. Chan, K. M. Tsui, H. C. Wu, Y. Hou, Y. C. Wu, and F. Wu, "Load/price forecasting and managing demand response for smart grid: Methodologies and challenges," *Signal Processing Magazine, IEEE*, vol. 29, no. 5, pp. 68–85, 2012.
- [34] J. W. Taylor, "Short-term load forecasting with exponentially weighted methods," *Power Systems, IEEE Transactions on*, vol. 27, no. 1, pp. 458–464, 2012.
- [35] M. Hagan and S. Behr, "The time series approach to short term load forecasting," *Power Systems, IEEE Transactions on*, vol. 2, no. 3, pp. 785–791, 1987.
- [36] M. H. Amini, O. Karabasoglu, M. D. Ilic, K. Boroojeni, and S. S. Iyengar, "Arima-based demand forecasting method considering probabilistic model of electric vehicles' parking lots," in *2016 IEEE Power and Energy Society General Meeting*. IEEE, 2015, pp. 1–5.
- [37] J. Contreras, R. Espinola, F. Nogales, and A. Conejo, "Arima models to predict next-day electricity prices," *Smart Grid, IEEE Transactions on*, vol. 18, no. 3, pp. 1014–1020, 2003.
- [38] M. Amini, A. Kargarian, and O. Karabasoglu, "Arima-based decoupled time series forecasting of electric vehicle charging demand for stochastic power system operation," *Electric Power Systems Research*, pp. 378–390, 2016.
- [39] G. Chicco, "Overview and performance assessment of the clustering methods for electrical load pattern grouping," *Energy*, vol. 42, no. 1, pp. 68–80, 2012.
- [40] J. D. Hamilton, *Time series analysis*. Princeton university press, 1994.
- [41] J. Friedman, T. Hastie, and R. Tibshirani, *The elements of statistical learning*. Springer, 2001.
- [42] N. Rao, R. Nowak, C. Cox, and T. Rogers, "Classification with the sparse group lasso," *Signal Processing, IEEE Transactions on*, vol. 64, no. 2, pp. 448–463, 2016.
- [43] A. Jung, G. Hannak, and N. Goertz, "Graphical lasso based model selection for time series," *Signal Processing Letters, IEEE*, vol. 22, no. 10, pp. 1781–1785, 2015.
- [44] D. Angelosante, J. A. Bazerque, and G. B. Giannakis, "Online adaptive estimation of sparse signals: Where rls meets the-norm," *Signal Processing, IEEE Transactions on*, vol. 58, no. 7, pp. 3436–3447, 2010.
- [45] N. Ludwig, S. Feuerriegel, and D. Neumann, "Putting big data analytics to work: Feature selection for forecasting electricity prices using the lasso and random forests," *Journal of Decision Systems*, vol. 24, no. 1, pp. 19–36, 2015.
- [46] Y. Goude, R. Nedellec, and N. Kong, "Local short and middle term electricity load forecasting with semi-parametric additive models," *Smart Grid, IEEE Transactions on*, vol. 5, no. 1, pp. 440–446, 2014.
- [47] R. Lockhart, J. Taylor, R. J. Tibshirani, and R. Tibshirani, "A significance test for the lasso," *Annals of statistics*, vol. 42, no. 2, p. 413, 2014.
- [48] R. J. Tibshirani, "The lasso problem and uniqueness," *Electronic Journal of Statistics*, vol. 7, pp. 1456–1490, 2013.
- [49] M. G. G'Sell, S. Wager, A. Chouldechova, and R. Tibshirani, "Sequential selection procedures and false discovery rate control," arXiv preprint arXiv:1309.5352, 2013.
- [50] A. Buja and L. Brown, "Discussion: a significance test for the lasso," *The Annals of Statistics*, vol. 42, no. 2, pp. 509–517, 2014.
- [51] National Oceanic and Atmospheric Administration, "National Centers for Environmental Information," 2015. [Online]. Available: <http://cdo.ncdc.noaa.gov>
- [52] K. G. Boroojeni, M. H. Amini, S. Bahrami, S. S. Iyengar, A. I. Sarwat, and O. Karabasoglu, "A novel multi-time-scale modeling for electric power demand forecasting: from short-term to medium-term horizon," *Electric Power Systems Research*, vol. 142, pp. 58–73, 2017.
- [53] A. Liaw and M. Wiener, "Classification and regression by randomforest," *R news*, vol. 2, no. 3, pp. 18–22, 2002.



Pan Li (S'16) received the B.S. degree in electrical engineering and the M.S. degree in systems engineering from Xi'an Jiaotong University, Xi'an, China, in 2011 and 2014, respectively. She also holds the engineering diploma from Ecole Centrale de Lille, France.

She is currently pursuing the Ph.D. degree at University of Washington, Seattle, WA, USA. Her current research interests include machine learning applications in demand response programs and in distribution networks.



Baosen Zhang received his Bachelor of Applied Science in Engineering Science degree from the University of Toronto in 2008; and his PhD degree in Electrical Engineering and Computer Sciences from University of California, Berkeley in 2013.

He was a Postdoctoral Scholar at Stanford University, affiliated with the Civil and Environmental Engineering and Management & Science Engineering. He is currently an Assistant Professor in Electrical Engineering at the University of Washington, Seattle, WA. His research interests are in power systems and cyberphysical systems. He was selected as one of Forbe's 30 under 30 in energy in 2015.



Yang Weng received the B.E. degree in electrical engineering from Huazhong University of Science and Technology, Wuhan, China; the M.Sc. degree in statistics from the University of Illinois at Chicago, Chicago, IL, USA; M.Sc. degree in machine learning of computer science, M.E. and Ph.D. degrees in electrical and computer engineering from Carnegie Mellon University (CMU), Pittsburgh, PA, USA.

After finishing his Ph.D., he joined Stanford University, Stanford, CA, USA, as the TomKat fellow for Sustainable Energy. He is currently an Assistant professor of electrical, computer and energy engineering at Arizona State University (ASU), Tempe, AZ, USA. His research interest is in the interdisciplinary area of power systems, machine learning, and renewable integration.

Dr. Weng received the CMU Dean's Graduate fellowship in 2010, the Best Paper Award at the International Conference on Smart Grid Communication (SGC) in 2012, the first ranking paper of SGC in 2013, Best Papers at the Power and Energy Society General Meeting in 2014, ABB fellowship in 2014, Golden Best Paper Award at the International Conference on Probabilistic Methods Applied to Power Systems in 2016.



Ram Rajagopal is an Assistant Professor of Civil and Environmental Engineering and Electrical Engineering (by Courtesy) at Stanford University, where he directs the Stanford Sustainable Systems Lab (S3L), focused on large scale monitoring, data analytics and stochastic control for infrastructure networks, in particular energy and transportation. His current research interests in power systems are in integration of renewables, smart distribution systems and demand-side data analytics. He has also built wireless sensing systems and large-scale analytics

for transportation and power networks. He holds a Ph.D. in Electrical Engineering and Computer Sciences and an M.A. in Statistics, both from the University of California Berkeley. He is a recipient of the Powell Foundation Fellowship, Berkeley Regents Fellowship, the NSF CAREER Award and the Makhoul Conjecture Challenge award. He holds more than 30 patents from his work, and has advised or founded various companies in the fields of sensor networks, power systems and data analytics.