

PIDS: An Intelligent Electric Power Management Platform

Yongqing Zheng,^{1,2,3} Han Yu,^{3,4*} Yuliang Shi,^{1,2,3*} Kun Zhang,^{1,2,5} Shuai Zhen,²
Lizhen Cui,^{1,3} Cyril Leung,^{3,6} Chunyan Miao^{3,4*}

¹School of Software, Shandong University (SDU), Jinan, China

²Dareway Software Co. Ltd, Jinan, China

³Joint SDU-NTU Centre for Artificial Intelligence Research (C-FAIR), Shandong University (SDU), Jinan, China

⁴School of Computer Science and Engineering, Nanyang Technological University (NTU), Singapore

⁵School of Information Science and Engineering, University of Jinan, Jinan, China

⁶Department of Electrical and Computer Engineering, The University of British Columbia (UBC), Vancouver, BC, Canada

*Corresponding Authors: {han.yu, ascymiao}@ntu.edu.sg, shiyuliang@sdu.edu.cn

Abstract

Electricity information tracking systems are increasingly being adopted across China. Such systems can collect real-time power consumption data from users, and provide opportunities for artificial intelligence (AI) to help power companies and authorities make optimal demand-side management decisions. In this paper, we discuss power utilization improvement in Shandong Province, China with a deployed AI application - the Power Intelligent Decision Support (PIDS) platform. Based on improved short-term power consumption gap prediction, PIDS uses an optimal power adjustment plan which enables fine-grained Demand Response (DR) and Orderly Power Utilization (OPU) recommendations to ensure stable operation while minimizing power disruptions and improving fair treatment of participating companies. Deployed in August 2018, the platform is helping over 400 companies optimize their power consumption through DR while dynamically managing the OPU process for around 10,000 companies. Compared to the previous system, power outage under PIDS through planned shutdown has been reduced from 16% to 0.56%, resulting in significant gains in economic activities.

Introduction

With the rapid development of the Chinese economy, the electricity demand by various industries is also increasing. Many provinces in China primarily rely on coal to generate electricity. This results in much environmental pollution; in addition, the coal supply limits the electric power generation (Meng et al. 2019). The inconsistent quality of coal results in high maintenance cost at the power plants. Seasonal climate changes and natural disasters may also cause short-term imbalances in the demand and supply of electricity, putting transient pressure on the power grid (Li, Pizer, and Wu 2019). These factors can result in reduced power generation capacities, fluctuations in electricity prices and sometimes forced power outages, thereby negatively affecting the stability and safe operation of the power grid.

At present, power grid management is carried out using two main approaches (Colak et al. 2016):

1. Supply side, through power grid scheduling;

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

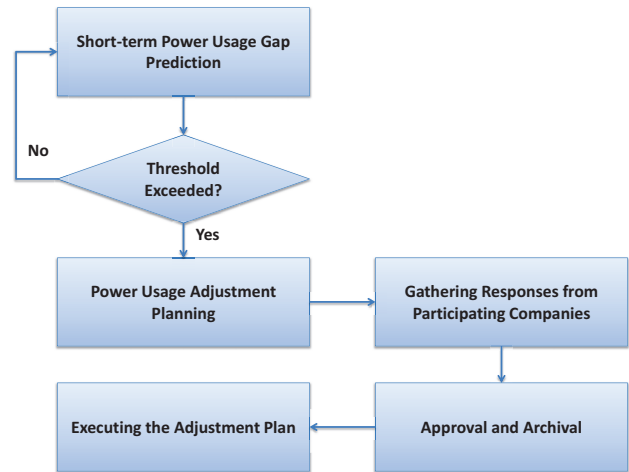


Figure 1: The general workflow of existing demand-side power management systems in China.

2. Demand side, through consumer usage scheduling.

In China, power generation capacity is still limited by coal fired power plants. There is not much room for optimization on the supply side. Thus, the focus of existing power management systems in China is generally on the demand side to mitigate peaks and troughs in power demand in order to maintain safe operation (Zhou and Yang 2015).

Existing demand-side power management systems in China generally operate following the steps below (Figure 1):

1. Short-term power consumption gap prediction: this step is typically performed by experts based on domain knowledge. In the case of Shandong province, the 3-day moving average values of previous power demand and supply are used as the basis for predicting the power consumption gaps in five days' time. If the predicted usage gap is larger than a predefined threshold value, the power consumption adjustment operations will be triggered.
2. Developing a hierarchical power consumption adjustment plan: based on the predicted power consumption

gap, the experts will decide on a specific period of time for power consumption adjustment, and divide the total amount of over-supply or shortage of power among the cities and counties in a province in a hierarchical fashion according to pre-defined rules. The adjusted power consumption quota allocated to each city/county will, in turn, be divided among companies located in the region which have signed agreements to participate in such operations.

3. Gathering responses to the plan from participating companies: the plan will then be forwarded to each of the eligible companies which are selected to participate in this round of power consumption adjustment for confirmation. If some of these companies prefer not to participate in this round (e.g., due to production scheduling conflicts), the plan will be manually adjusted to the extent allowed by the current situation for another round of confirmation.
4. Submitting the plan for approval: the finalized plan is then submitted for approval by the provincial power management authority and archived.
5. Executing the plan: execute the approved plan for power adjustment, either in the form of Demand Response (DR) (Strasser et al. 2015) through increase/reduction in electricity pricing, or Orderly Power Utilization (OPU) (Srivastava, Pandey, and Singh 2016) through planned shutdown during a given period for the selected companies.

From the above description, it can be observed that Steps 1 and 2, which are key to the performance of a power consumption management system, are currently performed manually by domain experts with computerized tools. Not only is this approach inefficient and unable to handle complex power consumption adjustment situations, it is also prone to human errors which might lead to accidents in production. In addition, the usage gap prediction approach currently adopted by such systems is lacking in terms of accuracy. Inaccurate predictions, in this case, introduces additional errors into the system which further affect its performance. Thus, the previously used system was unable to achieve rapid response to fluctuations in power demand and supply while minimizing disruption to economic activities.

In order to address this important limitation, we developed an artificial intelligence (AI) empowered dynamic power consumption management system - the Power Intelligent Decision Support (PIDS) platform. A novel short-term load forecasting model based on Wavelet Decomposition and Long Short-Term Memory (WD-LSTM) is incorporated into its AI Engine. It combines influencing factor analysis, wavelet decomposition feature extraction, triple order exponential smoothing (Holt-Winters) time series analysis, and Long Short-Term Memory (LSTM) networks to improve power consumption gap prediction. Based on the improved prediction results, PIDS computes an optimal power consumption adjustment plan which enables fine-grain adjustment of power consumption through joint objective constraint optimization to ensure safe operation while minimizing power disruptions and providing fair treatment of participating companies (Yu et al. 2019b; Zheng et al. 2019).

The platform provides detailed analysis for transparency in decision support.

PIDS has been deployed throughout Shandong Province since August 2018. It has significantly improved short-term power consumption prediction accuracy compared to the previous approach used in the province, and is helping over 400 companies optimize their power consumption through DR while dynamically managing the OPU process for around 10,000 companies. Compared to the previous system, planned shutdown under PIDS has been reduced from 16% to 0.56%, resulting in significantly reduced disruption to economic activities.

Application Description

In this section, we provide a detailed description of the system design of PIDS. The system architecture of PIDS is shown in Figure 2. It comprises of four tiers:

1. The Data Collection tier aggregates related data from multiple sources. These include load data and user profile information from the Electricity Information Collection System, meteorological data from the China Meteorological Data Service Center, and holiday data from public holiday calendars.
2. The Modeling tier which performs short-term load forecasting. Load forecasting is the basis for subsequent power consumption adjustment. Its accuracy will affect the effectiveness of subsequent steps. The proposed WD-LSTM model can meet the accuracy needed for our purpose.
3. The Utility tier includes DR and OPU. Based on the results of short-term load forecasting, this tier computes the

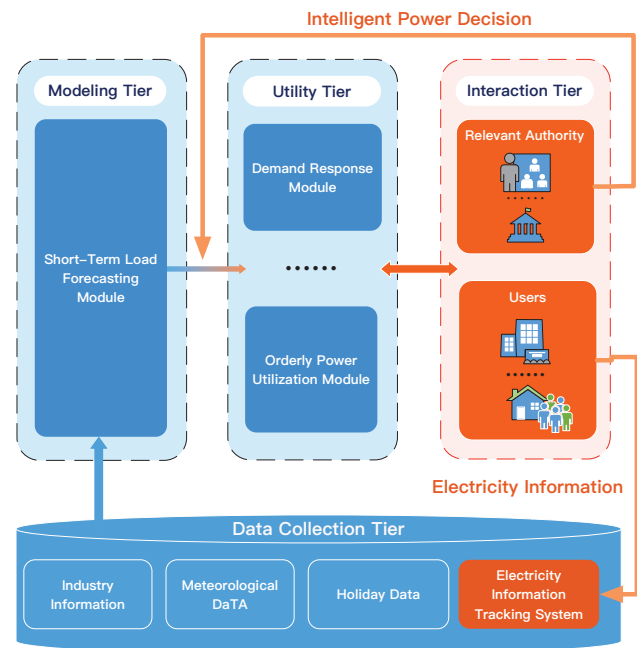


Figure 2: The architecture of the PIDS platform.

optimal power consumption adjustment quota for participating companies and sends out DR invitations to them. It is also responsible for recommending OPU operations to the authority in case DR alone is not enough to ensure safety.

- In the Interaction tier, users can query their historical power consumption behavior analysis reports. They can also accept/reject invitations to the DR operations. The confirmed power consumption adjustment amounts are delivered to the Electricity Information Collection System as feedbacks. The relevant authority can audit the processes of OPU and DR through the system with human interpretable explanations generated by the AI Engine, and modify future policies based on these results.

Use of AI Technology

The use of AI technologies in PIDS mainly focuses on performing two key steps in power consumption management: 1) predicting the short-term power consumption gap, and 2) optimizing the selection of companies to join DR and OPU. In this section, we describe these two parts in detail.

Short-Term Load Forecasting

One of the key functionalities of PIDS which relies heavily on AI techniques is predicting the short-term usage gap of power consumption. Accurate prediction of this value is important for the subsequent DR and OPU operations, in order to ensure safe operation of the power grid. Power consumption management in China generally faces the following challenges:

- Missing data: there are many possible causes for this problem. For example, hardware or software faults during data collection, the data collection mechanism cannot keep up with the speed of data being generated during peak usage periods, and high cost of data collection. A method for complementing the missing data is thus required.
- Noisy data: the collection of massive amounts of power consumption data over time is affected by random noise and environmental conditions. This negatively impacts the analysis of factors influencing the power load when data are viewed as a time series. Denoising is thus required during feature extraction.
- External influencing factors: power consumption can be affected by many factors external to the power grid (e.g., temperature which affects heating and cooling needs, public holidays which affect power needs of certain geographic locations and industries, and natural disasters). The accuracy of short-term power consumption gap prediction can be significantly improved by explicitly taking such factors into account.

To address these challenges, a novel short-term load forecasting model - WD-LSTM - is incorporated into the AI Engine of the PIDS platform. The conceptual framework of WD-LSTM is shown in Figure 3. It combines influencing factor analysis, wavelet decomposition feature extraction, triple order exponential smoothing (Holt-Winters) time

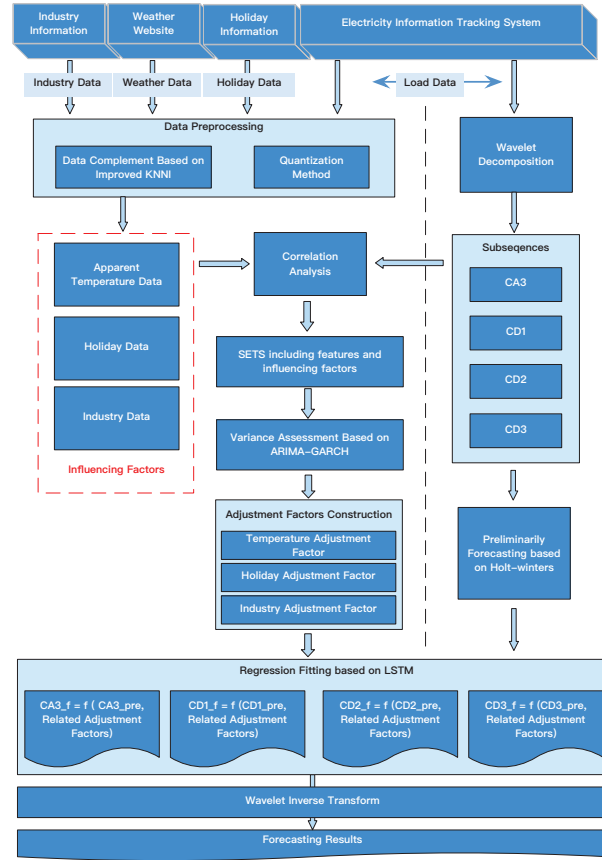


Figure 3: The conceptual framework of WD-LSTM.

series analysis and Long Short-Term Memory (LSTM) networks. The model uses wavelet decomposition to extract the main features of load data, analyzes its correlation with temperature, holidays and industry influencing factors, and then constructs corresponding adjustment factors.

To deal with the problem of noisy power consumption data, a 3-layer wavelet decomposition and reconstruction method is used in WD-LSTM. The four resulting subsequences and data concerning temperature, public holidays and industry specific information are used to perform correlation analysis to obtain the set of power load features and related influencing factors. Then, for each influencing factor, variance assessment based on the ARIMA-GARCH model (Tan et al. 2010) is performed to compute the adjustment values. The preliminary forecast for each feature subsequence is obtained using the Holt-Winters algorithm (Gelper, Fried, and Croux 2010). Finally, the forecasting result and the adjustment values are used as the input to the LSTM network (Gers, Schmidhuber, and Cummins 1999) to perform regression forecasting and wavelet inverse transformation to obtain the best forecasting results. More details of WD-LSTM can be founded in (Liu et al. 2019; Zheng et al. 2019).

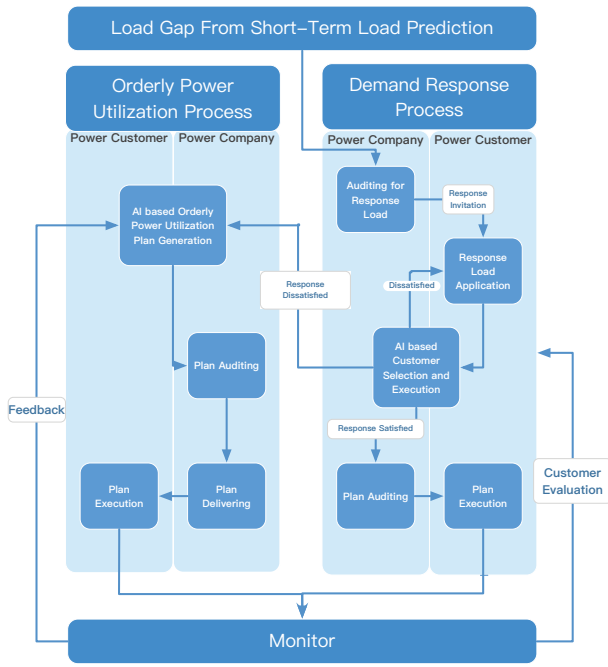


Figure 4: The architecture of the PIDS platform.

Dynamic Power Consumption Adjustment

In PIDS, there are a set of N companies which have signed contracts with the provincial government agreeing to participate in power consumption adjustment operations in the following year in exchange for preferential electricity rates. The actual amount of concessions received will depend on the level of participation by these companies.

Power consumption adjustments can be divided into two levels: 1) demand response (DR) and 2) orderly power utilization (OPU) as shown in Figure 4. DR can be used to reduce power demands during peak periods, or increase power demands during trough periods. Under DR, a user reduces (or increases) its power consumption by an agreed amount over a specific period of time (e.g., through partially shutting down operations or boosting production activities) in exchange for a lower electricity price during the DR period. During peak periods, if the amount of power consumption reduction in a given round of DR operation is not enough to bridge the power demand-supply gap, OPU will be triggered. PIDS will then select some target companies to be forcefully powered down during the specific period in order to ensure safe operation of the power grid. The selection of companies for DR and OPU must not only satisfy safety constraints, but also minimize economic loss and ensure fair treatment of the participating companies.

Demand Response (DR) In order to achieve these objectives, we have extended the framework in (Yu et al. 2013a; 2015; 2016; 2017) which has been deployed in our social insurance service provision platform (Zheng et al. 2018) to enable PIDS to dynamically allocate power consumption adjustment quotas among participating companies.

Depending on which industry a company belongs to, the economic output per unit time (e.g., in an hour) of one company may be different from that of another company. In PIDS, this value for company i is denoted as v_i and normalized to the range $[0, 1]$. In the contract for participating in dynamic power consumption adjustment, a company i specifies the maximum amount of power reduction, x_i^{\max} , it is willing to tolerate in any given round of DR operation.

From the system perspective, the sacrifice, $Y_i(t)$, made by each company i during round t can be modelled by as a queuing system:

$$Y_i(t+1) = \max[Y_i(t) + x_i(t) - 1_{[x_i(t)=0]}\mu_i(t), 0] \quad (1)$$

where $x_i(t)$ is the power reduction quota allocated to i in round t . $\mu_i(t)$ is the actual amount of electricity i consumes in round t . However, $\mu_i(t)$ is only counted if i is not required to reduce power consumption through DR or OPU in round t as indicated by the function $1_{[x_i(t)=0]}$ (which equals to 1 if only if $x_i(t) = 0$; and equals to 0 otherwise).

Based on (Yu et al. 2017), we model the distribution of sacrifice among the companies in round t with a Lyapunov function, $L(t)$, defined as:

$$L(t) = \frac{1}{2} \sum_{i=1}^N Y_i^2(t) \quad (2)$$

which is, in essence, the l_2 -norm value. A large $L(t)$ indicates that either a few companies are making huge sacrifices, or a number of companies are making large sacrifices. Both of these situations shall be avoided as much as possible. The fluctuation in company sacrifices over T rounds can be expressed as:

$$\Delta = \frac{1}{T} \sum_{t=0}^{T-1} [L(t+1) - L(t)]. \quad (3)$$

Thus, minimizing Δ achieves two dimensions of fair treatment: 1) the distribution of sacrifice among companies within each round is even over time, and 2) the fluctuations in the sacrifices made by the companies over time remain small.

PIDS also takes the economic impact of the companies into account when prioritizing which ones shall join DR in a given round. For this purpose, we model the normalized loss incurred by a company i reducing power consumption by $x_i(t)$ as $v_i x_i(t)$. Thus, the time-averaged economic loss incurred by power consumption adjustment in the system is:

$$F = \frac{1}{T} \sum_{t=0}^{T-1} \sum_{i=1}^N v_i x_i(t). \quad (4)$$

In order to minimize economic loss and ensure fair treatment of participating companies while achieving safe operation, we need to minimize $\rho F + \Delta$ subject to certain constraints. Here, $\rho > 0$ is a weight value allowing the system administrator to indicate the relative importance between minimizing economic loss and ensuring fair treatment of participating companies. Thus, we have:

Minimize:

$$\frac{1}{T} \sum_{t=0}^{T-1} \sum_{i=1}^N x_i(t) [\rho v_i + Y_i(t)] \quad (5)$$

Subject to:

$$x_i(t) \leq x_i^{\max}, \forall i, t \quad (6)$$

Minimizing equation (5) jointly minimizes the time-averaged expected disruption to economic activities and the sacrifices made by participating companies. Constraint (6) ensures that during DR operations, any company i will not be asked to adjust their power consumption by more than x_i^{\max} . For simplicity of discussion, we denote the term $[\rho v_i + Y_i(t)]$ as the *DR suitability index*, $\Theta_i(t)$. In the PIDS platform, the system administrators have decided to give equal emphasis on minimizing economic loss and ensuring fair treatment of participating companies. Thus, the value of variable ρ is set in such a way that the terms ρv_i and $Y_i(t)$ from equation (5) are of the same order of magnitude.

Algorithm 1 PIDS Demand Response

Require: The predicted power consumption gap at some future round t , $G(t)$; $Y_i(t)$, $\forall i, t$; v_i , $\forall i$; and ρ .

- 1: Sort all i in ascending order of their $\Theta_i(t)$ values;
- 2: **for** each i **do**
- 3: **if** $G(t) \geq x_i^{\max}$ **then**
- 4: $x_i(t) = x_i^{\max}$;
- 5: **else**
- 6: $x_i(t) = G(t)$;
- 7: **end if**
- 8: $G(t) \leftarrow G(t) - x_i(t)$;
- 9: **end for**
- 10: **return** $\{x_1(t), x_2(t), \dots, x_N(t)\}$;

The algorithm for PIDS to allocate power consumption adjustment quotas to participating companies is shown in Algorithm 1. It is only triggered if the predicted power consumption gap $G(t)$ is larger than a predefined threshold value. The required inputs can all be tracked and stored by the PIDS platform. The algorithm has $O(N \log N)$ time complexity if mergesort (Katajainen and Träff 1997) is adopted, making it suitable for large-scale operations.

The PIDS platform coordinates the allocation of power consumption adjustment quotas by sending messages to selected companies and gathering their responses. Intuitively, if a power disruption to a company i incurs low economic cost, and i has not made much sacrifice by participating in recent rounds of DR operations, the priority for i to join this round of DR operation is increased; and vice versa. If some companies decline this round of invitation, PIDS re-runs Algorithm 1, with all the companies which have accepted or declined the invitations excluded from the eligible solution set, to compute a new solution for DR. Since the authority only informs the AI Engine of high level preferences through setting key DR parameters without insisting on approving each AI recommendation, the nature of AI operations in this part is *human-over-the-loop* (Yu et al. 2019a).

Intelligent Orderly Power Utilization (OPU) If Algorithm 1 cannot find a feasible solution to bridge the predicted gap between power demand and supply, PIDS will trigger the OPU operation. It generates a preliminary proposal for selected companies to completely shut down their operations during the predicted power consumption gap period based on the load characteristics of the companies which have signed contracts with the provincial authority to agree to participate in OPU. This proposal is submitted to the Commission of Economy and Information Technology for approval.

The administrators from the Commission of Economy and Information Technology logs into the PIDS platform, and enters the government department management interface to review the OPU plan. Since the recommendations for performing OPU are made based on only predicted power consumption gaps, the authority needs to make the final decision on whether to approve such recommendations. Thus, the nature of AI operations in this part is *human-in-the-loop* (Yu et al. 2019a). If the OPU plan is approved, the execution of the OPU plan will be managed by PIDS.

Application Use and Payoff

PIDS has been deployed across Shandong Province, China since August 2018. The province consists of 16 prefectures, 140 counties and 1,941 townships with a total population of around 100 million¹. It has been used by the provincial authority to manage the demand side electricity usage by industrial users. At the time of submission of this paper, over 400 companies signed agreements to participate in DR operations, and over 10,000 companies have signed agreements to participate in OPU operations through PIDS. In this section, we discuss the impact of the PIDS platform. We compare the performance achieved by PIDS with data from the previous power consumption management system used by Shandong Province during the period of August 2017 to July 2018. The performance data of PIDS were gathered from August 2018 to July 2019.

With WD-LSTM, the PIDS platform achieved a root mean square error (RMSE) of less than 2.5% when predicting the short-term power consumption gap during the one year period of deployment. This represents a significant 86.40% reduction compared to the RMSE of 18.38% achieved by the 3-day moving averaged-based prediction approach adopted by the previous power consumption management system in Shandong.

The first round of province-wide DR operation occurred in August 2018. A total of 264 companies received DR invitations to bridge the power consumption gap of 555,800kW. Eventually, 201 companies accepted the invitations and adjusted their power consumption by 439,000kW through DR (accounting for 79% of the power gap). The remaining power gap was bridged through OPU.

Another round of province-wide DR operation occurred in December 2018. A total of 188 companies boosted their power consumption at noon time (12:00 pm to 13:00 pm) to bring the power demand trough up by 291,900kW. Then, on

¹<https://en.wikipedia.org/wiki/Shandong>

the same day, 327 companies reduced their power consumption from 17:00 pm to 18:00 pm to bring the power demand peak down by 640,900kW. This round of power consumption adjustment reduced the difference between the peak and trough of the power demand by 15.74% (from 5,566,400 kW to 4,690,200 kW), thereby significantly improving the power demand curve.

With the help of PIDS, an average of 76% of participating companies did not experience significant disruptions to their power consumption on the day of power consumption adjustment. This is an improvement of more than 58% compared to the average of only 48% of participating companies which managed to achieve this under the previous system.

With the help of PIDS, participating companies only had to reduce their power consumption by an average of 0.56% on the day of power consumption adjustment. In contrast, under the previous system, the corresponding figure was 16%.

Under PIDS, within the week after the power consumption adjustment operation, participating companies can recoup the reduction in power consumption safely, thereby minimizing the negative impact on economic activities. In contrast, an average of 8% of production output was permanently lost due to power consumption adjustment under the previous system. Overall, PIDS has achieved much a better performance compared to the previous system over the one year period it was deployed.

Application Development and Deployment

The PIDS platform was developed using Java and JSP Programming Languages by Dareway Software Co Ltd.², Jinan, Shandong, China. The platform is built on Hadoop for storing the big data in power consumption. When developing the AI Engine, we have evaluated five well-established potential models for short-term power consumption gap prediction:

1. Support Vector Regression (SVR) (Drucker et al. 1996): this is a special version of the Support Vector Machine (SVM) designed for performing regression.
2. AutoRegressive Integrated Moving Average (ARIMA) (Makridakis and Hibon 1997): this is a model that captures a suite of different standard temporal structures in time series data. It can be applied in cases where data are non-stationary. An initial differencing step is carried out once or multiple times to deal with such non-stationarity.
3. Long Short-Term Memory (LSTM) neural network (Gers, Schmidhuber, and Cummins 1999): this is a deep recurrent neural network (RNN). In contrast to standard feed-forward neural networks, LSTM has feedback connections. It not only can process single data points, but also sequences of data (e.g., power consumption over time).
4. Holt-Winters Forecasting (Gelper, Fried, and Croux 2010): this method is also known as triple exponential smoothing, which applies exponential smoothing three

²<http://www.dareway.com.cn/>

Table 1: Power consumption prediction test results.

Algorithm	Performance	RMSE	Reduction in RMSE by WD-LSTM
WD-LSTM		9.76%	-
ARIMA		24.42%	-60.03%
Holt-Winters		21.89%	-55.41%
SVR		21.26%	-54.09%
LSTM		14.68%	-33.51%

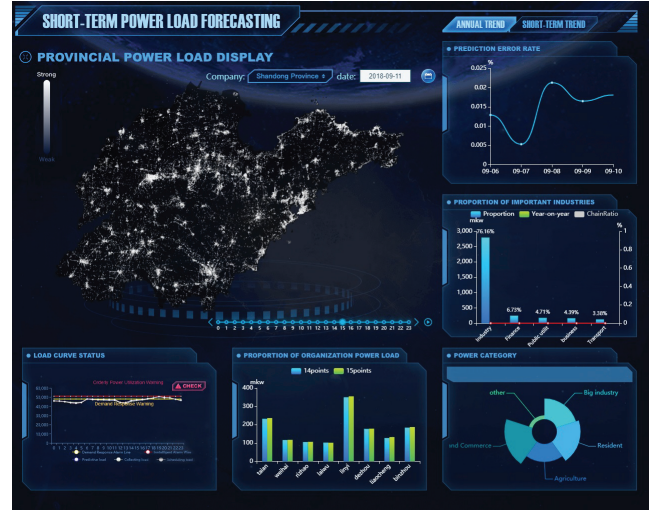


Figure 5: The PIDS user interface showing the alert on predicted power consumption gap and related analysis results (Zheng et al. 2019).

times. It is commonly used when there are three high frequency signals to be removed from a time series under study.

5. WD-LSTM (Liu et al. 2019; Zheng et al. 2019): this is our proposed approach for short-term power consumption gap prediction designed to address the specific challenges facing Chinese power consumption management systems, which was described in the previous section.

The five candidate approaches have been tested on a power demand and supply dataset³ collected by the previously used power consumption management system from a city in Shandong Province, China. The dataset spans a one year period, across different seasons. The task for the candidate approaches is to predict the power demand and supply five days in the future. The RMSEs achieved by the candidate approaches are listed in Table 1. The results showed that WD-LSTM achieves the lowest RMSE at 9.76% based on the historical test dataset. Compared to the best alternative approach (i.e. LSTM), the RMSE achieved by WD-LSTM is lower by 33.51%. Therefore, WD-LSTM has been selected to be implemented in the AI Engine of PIDS.

Figure 5 shows the user interface through which PIDS dis-

³The dataset is available upon request via email to the contact authors as the requesters will be asked to agree on certain terms of use.

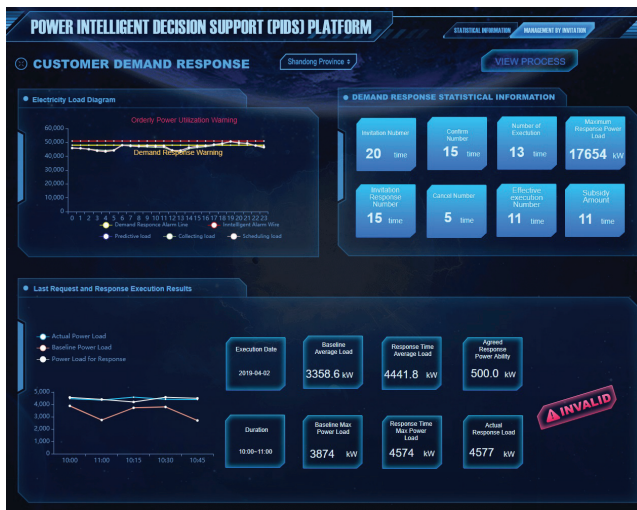


Figure 6: The PIDS user interface showing decision support functions related to demand response optimization.

plays the predicted power consumption gap and related analysis for the administrators of the power grid in Shandong Province. The actual deployed system uses the Chinese language. This is a translated version for readers who do not speak Chinese. In the example, on 11-Sep-2018, the WD-LSTM model predicted a near-term power consumption gap on 15-Sep-2018 that would exceed the pre-defined safe operation parameters. It issued an alert in the Load Curve Status panel. Detailed breakdowns of the year-to-date power consumption by different types of customers are shown in different panels. The historical prediction error rates are also plotted in the Prediction Error Rate panel on the top-right hand corner of the screen. The additional analysis is shown to provide transparency and help the administrators make informed decisions on whether to act on the alerts.

Figure 6 shows the user interface through which PIDS displays recommendations for selected companies to participate in a round of demand response operation, as well as a summary of compliance by these companies on 02-Apr-2019. In this instance, due to poor compliance from the selected companies for DR, forced power shutdown through OPU was activated. A video demonstration of the PIDS platform can be accessed online⁴.

Maintenance

As time goes by, workflows, personnel and operating parameters in PIDS may change. Because of the separation of concerns through modular system design, such updates can be performed without affecting the AI Engine. The AI algorithms did not need to be modified since deployment in August 2018. At the point of submission of this paper, there has not been any AI maintenance task. At the end of August 2019, after one year of operation, a model review for WD-LSTM has been scheduled. This task will be carried out over

⁴<https://www.youtube.com/watch?v=3dcX1YDwT0Q>

the subsequent month to determine if the WD-LSTM model should be updated.

Conclusions and Future Work

In this paper, we reported on our experience using AI to address the challenges of dynamically managing industrial electric power consumption in Shandong Province, China. We developed the PIDS platform to provide data-driven intelligent power consumption adjustment decision support for the provincial authority. By improving the accuracy of the short-term power consumption gap prediction and dynamically optimizing the selection of companies to join demand response and orderly power utilization operations, PIDS provides fine-grain adjustments of the power demand curve in order to ensure safe operation while minimizing power disruptions and providing fair treatment of participating companies. Since deployment, PIDS has helped over 400 companies in Shandong Province optimize their power consumption through DR while dynamically managing the OPU process for around 10,000 companies. The platform has demonstrated significant advantages in terms of improving the management of the power grid with minimal impact on economic activities compared to the previous system. Further plans to deploy PIDS in other parts of China have been put in place. The experience gained is being analysed to help revise related policies in China.

In subsequent work, we will investigate how to incorporate Stackelberg game theory (von Stackelberg 2011) into PIDS to enable the dynamic pricing of electric power. In this way, the platform may be able to make use of price signals to better motivate companies to participate in power consumption adjustment operations. We are also looking into applying explainable AI approaches (Fan and Toni 2015; Zeng et al. 2019) to automatically generate explanations for the AI recommendations to help administrators better understand the reasoning processes in order to improve user acceptance and trust in AI (Yu et al. 2013b; 2018). In addition, as smart meter devices start to be incorporated into the power grid infrastructure in Shandong Province, massive amounts of usage behaviour data will be generated by companies and households. Not only is it expensive to transmit and store such data, but private information such as patterns in people's daily life might be inferred from such data as well. We will also investigate how to apply privacy-preserving machine learning techniques such as federated learning (Gao et al. 2019; Yang et al. 2019) in PIDS so as to enable collaboration across power companies in compliance with privacy-protection laws such as the General Data Protection Regulation (GDPR) (Voigt and Bussche 2017).

Acknowledgments

This research is supported by the Nanyang Assistant Professorship (NAP); AISG-GC-2019-003; NRF-NRFI05-2019-0002; NTU-SDU-CFAIR (NSC-2019-011); National Key Research and Development Plan of China (No. 2018YFB1003804); the Taishan Industrial Experts Program of Shandong Province of China (No. tscy20150305); the Key Research & Development Program of Shandong Province of China (No. 2016ZDJS01A09); the NSFC No.

91846205; the Innovation Method Fund of China No. 2018IM020200; SDNFSC No. ZR2017ZB0420.

References

- Colak, I.; Sagiroglu, S.; Fulli, G.; Yesilbudak, M.; and Catalin-FelixCovrig. 2016. A survey on the critical issues in smart grid technologies. *Renewable and Sustainable Energy Reviews* 54:396–405.
- Drucker, H.; Burges, C. C.; Kaufman, L.; Smola, A. J.; and Vapnik, V. N. 1996. Support vector regression machines. In *Advances in Neural Information Processing Systems (NIPS'96)*, 155–161.
- Fan, X., and Toni, F. 2015. On computing explanations in argumentation. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI-15)*, 1496–1502.
- Gao, D.; Liu, Y.; Huang, A.; Ju, C.; Yu, H.; and Yang, Q. 2019. Privacy-preserving heterogeneous federated transfer learning. In *Proceedings of the 2019 IEEE International Conference on Big Data (IEEE BigData'19)*.
- Gelper, S.; Fried, R.; and Croux, C. 2010. Robust forecasting with exponential and Holt–Winters smoothing. *Journal of Forecasting* 29(3):285–300.
- Gers, F. A.; Schmidhuber, J.; and Cummins, F. A. 1999. Learning to forget: continual prediction with LSTM. In *Proceedings of the 9th International Conference on Artificial Neural Networks (ICANN'99)*, 850–855.
- Katajainen, J., and Träff, J. L. 1997. A meticulous analysis of mergesort programs. In *Proceedings of the 3rd Italian Conference on Algorithms and Complexity (ICAC'97)*, 217–228.
- Li, Y.; Pizer, W. A.; and Wu, L. 2019. Climate change and residential electricity consumption in the Yangtze River Delta, China. *Proceedings of the National Academy of Sciences (PNAS)* 116(2):472–477.
- Liu, Y.; Zhang, K.; Zhen, S.; Guan, Y.; and Shi, Y. 2019. WRL: A combined model for short-term load forecasting. In *Proceedings of the 3rd International Joint Conference on Web and Big Data (APWeb-WAIM'19)*, 35–42.
- Makridakis, S., and Hibon, M. 1997. Arima models and the box-jenkins methodology. *Journal of Forecasting* 16(3):147–163.
- Meng, W.; Zhong, Q.; Chen, Y.; Shen, H.; Yun, X.; Smith, K. R.; Li, B.; Liu, J.; Wang, X.; Ma, J.; Cheng, H.; Zeng, E. Y.; Guan, D.; Russell, A. G.; and Tao, S. 2019. Energy and air pollution benefits of household fuel policies in northern China. *Proceedings of the National Academy of Sciences (PNAS)* 116(34):16773–16780.
- Srivastava, A. K.; Pandey, A. S.; and Singh, D. 2016. Short-term load forecasting methods: A review. In *International Conference on Emerging Trends in Electrical Electronics and Sustainable Energy Systems*.
- Strasser, T.; Andrén, F.; Kathan, J.; Cecati, C.; Buccella, C.; Siano, P.; Leitão, P.; Zhabelova, G.; Vyatkin, V.; and Vrba, P. 2015. A review of architectures and concepts for intelligence in future electric energy systems. *IEEE Transactions on Industrial Electronics* 62(4):2424–2438.
- Tan, Z.; Zhang, J.; Wang, J.; and Xua, J. 2010. Day-ahead electricity price forecasting using wavelet transform combined with ARIMA and GARCH models. *Applied Energy* 87(11):3606–3610.
- Voigt, P., and Bussche, A. v. d. 2017. *The EU General Data Protection Regulation (GDPR): A Practical Guide*. Springer Publishing Company, Incorporated, 1st edition.
- von Stackelberg, H. 2011. *Market Structure and Equilibrium*. Springer.
- Yang, Q.; Liu, Y.; Chen, T.; and Tong, Y. 2019. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology* 10(2):12:1–12:19.
- Yu, H.; Miao, C.; An, B.; Leung, C.; and Lesser, V. R. 2013a. A reputation management model for resource constrained trustee agents. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI'13)*, 418–424.
- Yu, H.; Shen, Z.; Leung, C.; Miao, C.; and Lesser, V. R. 2013b. A survey of multi-agent trust management systems. *IEEE Access* 1(1):35–50.
- Yu, H.; Miao, C.; Shen, Z.; Leung, C.; Chen, Y.; and Yang, Q. 2015. Efficient task sub-delegation for crowdsourcing. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI-15)*, 1305–1311.
- Yu, H.; Miao, C.; Leung, C.; Chen, Y.; Fauvel, S.; Lesser, V. R.; and Yang, Q. 2016. Mitigating herding in hierarchical crowdsourcing networks. *Scientific Reports* 6(4):doi:10.1038/s41598-016-0011-6.
- Yu, H.; Miao, C.; Chen, Y.; Fauvel, S.; Li, X.; and Lesser, V. R. 2017. Algorithmic management for improving collective productivity in crowdsourcing. *Scientific Reports* 7(12541):doi:10.1038/s41598-017-12757-x.
- Yu, H.; Shen, Z.; Miao, C.; Leung, C.; Lesser, V. R.; and Yang, Q. 2018. Building ethics into artificial intelligence. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI'18)*, 5527–5533.
- Yu, H.; Liu, Y.; Wei, X.; Zheng, C.; Chen, T.; Yang, Q.; and Peng, X. 2019a. Fair and explainable dynamic engagement of crowd workers. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI'19)*, 6575–6577.
- Yu, H.; Miao, C.; Zheng, Y.; Cui, L.; Fauvel, S.; and Leung, C. 2019b. Ethically aligned opportunistic scheduling for productive laziness. In *Proceedings of the 2nd AAAI/ACM Conference on AI, Ethics, and Society (AI/ES-19)*, 45–51.
- Zeng, Z.; Miao, C.; Leung, C.; Shen, Z.; and Chin, J. J. 2019. Computing argumentative explanations in bipolar argumentation frameworks. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI-19)*, 10079–10080.
- Zheng, Y.; Yu, H.; Cui, L.; Miao, C.; Leung, C.; and Yang, Q. 2018. SmartHS: An AI platform for improving government service provision. In *Proceedings of the 30th AAAI Conference on Innovative Applications of Artificial Intelligence (IAAI-18)*, 7704–7712.
- Zheng, Y.; Yu, H.; Zhang, K.; Shi, Y.; Leung, C.; and Miao, C. 2019. Intelligent decision support for improving power management. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI'19)*, 6587–6589.
- Zhou, K., and Yang, S. 2015. Demand side management in China: The context of China's power industry reform. *Renewable and Sustainable Energy Reviews* 47:954–965.