Energy Proceedings
Vol 42, 2024
ISSN 2004-2965

DAF-GAN: Day-Ahead Forecasting of Building HVAC Energy Consumption Using Multi-Channel Generative Adversarial Networks

Yichuan X. Ma*, Lawrence K. Yeung

Department of Electrical and Electronic Engineering, The University of Hong Kong, Hong Kong SAR, PRC

(*Corresponding Author: xma.hku@gmail.com)

ABSTRACT

In the pursuit of sustainability and energy efficiency, accurate short-term prediction of HVAC energy consumption is crucial. Deep learning emerges as a promising solution for handling diverse data challenges in building HVAC systems. While deep generative learning excels in computer vision, its potential in predicting energy consumption remains largely untapped. This study first introduces a novel framework, transforming forecasting into a conditional generative problem in the temporal domain. We then propose **DAF-GAN**, an image inpainting-based data-driven method for **D**ay-**A**head **F**orecasting of buildings' HVAC energy consumption using multi-channel **G**enerative **A**dversarial **N**etworks (GANs). In day-ahead forecasting tasks across eleven real-world buildings, DAF-GAN exhibits relative improvements of 17% to 68% across four different error metrics compared to six traditional and deep learning models. DAF-GAN also demonstrates less bias and superior stability when applied to different buildings, holding promise for enhancing energy-efficient building automation and management.

Keywords: building energy forecasting, HVAC, conditional generative adversarial network, deep learning, image inpainting, artificial intelligence

NOMENCLATURE

Abbreviations	
1DCNN	One-Dimensional Convolutional Neural Network
2DCNN	Two-Dimensional Convolutional Neural Network
ARIMA	Autoregressive Integrated Moving Average
CNN	Convolutional Neural Network
CV-RMSE	Coefficient of Variation of Root Mean Square Error
GAN	Generative Adversarial Network
GASF	Gramian Angular Summation Field
GRU	Gated Recurrent Unit
HVAC	Heating, Ventilation and Air-Conditioning
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
MBE	Mean Bias Error
MC-1DCNN	Multi-Channel One-Dimensional Convolutional Neural Network
MLP	Multilayer Perceptron
RMSE	Root Mean Square Error
Symbols	
y_k	Actual Data Point
$\widehat{y_k}$	Estimated Data Point

1. INTRODUCTION

In the quest for sustainability and energy efficiency, the effective management of Heating, Ventilation, and Air Conditioning (HVAC) systems within the built environment assumes a pivotal role. HVAC systems are the most energy-consuming service worldwide, contributing to 38% (12%) of global building (final) energy use [1], making their effective management a key driver in achieving energy conservation objectives. The cornerstone of this endeavour is the accurate and timely prediction of HVAC energy consumption, which offers vital insights for proactive energy management and cost mitigation.

Despite remarkable strides in HVAC system design and control strategies, the ability to forecast energy consumption at fine temporal resolutions, particularly for the short term, remains a formidable challenge. Accurate short-term predictions are pivotal for

This is a paper for 15th International Conference on Applied Energy (ICAE2023), Dec. 3-7, 2023, Doha, Qatar.

optimizing HVAC operations and enabling demand-side management. Traditional forecasting methodologies, such as Autoregressive Integrated Moving Average (ARIMA) [2], often fall short of capturing the intricate dynamics inherent in HVAC systems, particularly when dealing with diverse and multivariate data sources.

Deep generative learning methods, such as generative adversarial networks (GANs) [3], have been continually demonstrating remarkable advantages in computer vision, while their potential in building energy consumption prediction remains largely untapped. Although certain studies have recognized the applicability of GANs in predicting building energy consumption, their utilization has been predominantly confined to indirect tasks such as feature engineering [4] or data augmentation [5]. The direct application of GANs for energy consumption prediction has yet to be extensively explored.

Recognizing these research gaps, in this paper, we first introduce a novel framework to reformulate forecasting into a conditional generative problem. Then we propose **DAF-GAN**, an image inpainting-based data-driven method for **D**ay-**A**head **F**orecasting of buildings' HVAC energy consumption using multi-channel **GAN**s, which directly employs GANs for precise building energy consumption prediction. To rigorously assess the performance of DAF-GAN, we meticulously selected and implemented a comprehensive suite of benchmark models, encompassing ARIMA, Multilayer Perceptron (MLP) [6], Convolutional Neural Network (CNN) [7] w/ or w/o multi-channel inputs, Long Short-Term Memory (LSTM) [8], and Gated Recurrent Unit (GRU) [9]. We defined a comprehensive set of five evaluation metrics and compared the performance of DAF-GAN and the selected benchmark models in our case studies, which utilized actual HVAC energy consumption data from eleven realistic buildings.

The remainder of this paper is structured as follows: Section 2 provides the theoretical and methodological details of our proposed framework and DAF-GAN; Section 3 outlines the case study implementation, including data sources and experimental design; Section 4 presents the results and discussions; and Section 5 concludes the paper and highlights our future research extensions related to this study.

2. THEORY AND METHODOLOGY

2.1 GANs and Conditional GANs

A GAN consists of two neural networks, a generator and a discriminator, engaged in a competitive process. The generator generates synthetic data to mimic real data, while the discriminator tries to distinguish between real and synthetic data. As training progresses, the generator learns to produce increasingly realistic data, while the discriminator becomes adept at differentiating real from fake data. This adversarial training leads to the generator producing high-quality data samples.

Conditional GANs [10] extend the GAN concept by introducing conditional information. In conditional GANs, the generator takes additional input in the form of conditional data, such as images with corrupted regions that require restoration. This conditioning empowers the generator to produce data samples that align with specific conditions.

2.2 A novel framework to reformulate forecasting into a conditional generative problem

Building upon the fundamental insight that "**forecasting is about generating data of future**" [10], we introduce a novel framework that reformulates forecasting into a problem of conditional generation.

This framework leverages the Gramian angular summation field (GASF) technique [11], which affords a unique, bijective image encoding of time series data. By encoding historical energy time series into GASF and strategically padding unknown future values with zeros, we effectively reframe the *forecasting* task into a specific *conditional generative problem*, i.e., an image inpainting problem (akin to the example illustrated in Fig. 1).

Fig. 1. Example of an image inpainting problem

As shown in Fig. 1, image inpainting is a computer vision problem that requires special techniques to fill in or restore missing or damaged portions of an image, effectively "painting" over the gaps with contextually appropriate visual information. Conditional GANs are particularly well-suited for image inpainting tasks due to their ability to generate high-quality, contextually relevant image content. By conditioning the generator on both the available image context and the target inpainting location, conditional GANs excel in producing high-quality reconstructions, effectively "completing" missing regions with plausible information. This makes them a valuable tool for the reformulated forecasting task.

2.3 *DAF-GAN: Day-ahead forecasting of buildings' HVAC energy consumption using multi-channel GANs*

Fig. 2 depicts the inference stage of our proposed DAF-GAN method, which also shows how it reformulates a forecasting problem into an image inpainting problem.

The multi-channel inputs are three-fold, including padded historical energy inputs (1 channel), meteorological inputs (5 channels) and date information inputs (2 channels). Previous studies have consistently underscored the substantial impact of meteorological conditions [13] and date information [14] on building energy consumption. Therefore, the integration of multiple channels to account for these influential factors holds the promise of enhancing predictive performance.

DAF-GAN leverages the UNet [15] architecture as the forecaster (generator) to produce inpainted HVAC energy consumption maps. In parallel, a two-dimensional CNN (2DCNN) discriminator is employed to assess the authenticity of the inpainted maps.

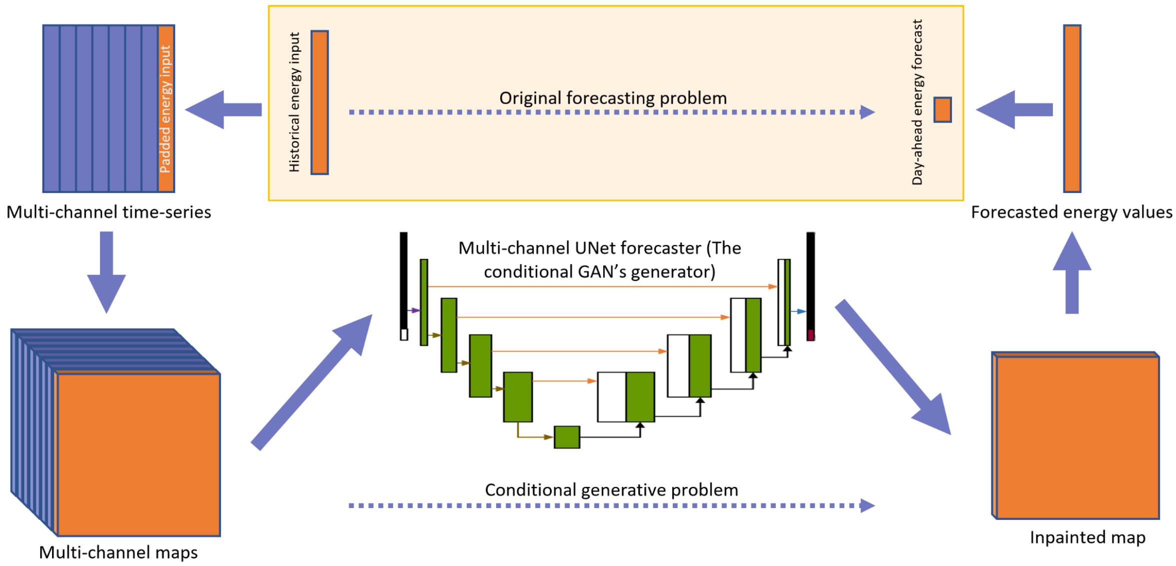

Fig. 2. Schematic representation of the proposed DAF-GAN method converting forecasting into a conditional generative problem

3. CASE STUDIES

3.1 *Data*

The building HVAC energy data and weather data used in this paper were drawn from the "Energy Detective" competition [16], including three years (2015-2017) of energy data collected from eleven office buildings in Shanghai. The holiday data were summarized based on the holiday schedules announced by the General Office of the State Council of the PRC.

3.2 *Experimental design*

3.2.1 Data splitting strategy

We utilize data from 2015-2016 as the training set (66.7%) and evenly split the data from 2017 into a validation set (16.7%) and a testing set (16.7%). This approach ensures a fair evaluation of models' performance on unseen data while using historical data for training. A tailored data cleansing procedure was applied before feeding the data to the models.

3.2.2 Evaluation metrics

To provide an extensive evaluation of the models' performance, five evaluation metrics are used to evaluate the performance of different models. These metrics include four error metrics, namely the mean absolute percentage error (MAPE), coefficient of variation of root mean square error (CV-RMSE or CV_{RMSE}), mean absolute error (MAE), root mean square error (RMSE), and one bias metric, namely the mean bias error (MBE). Definitions of the error and bias metrics are shown below.

$$MAPE = \frac{1}{n}\sum_{k=1}^{n}\left|\frac{y_k - \widehat{y_k}}{y_k}\right|$$

$$CV_{RMSE} = \sqrt{\frac{1}{n}\sum_{k=1}^{n}(y_k - \widehat{y}_k)^2} \Big/ \left(\frac{1}{n}\sum_{k=1}^{n} y_k\right)$$

$$MAE = \frac{1}{n}\sum_{k=1}^{n}|y_k - \widehat{y_k}|$$

$$RMSE = \sqrt{\frac{1}{n}\sum_{k=1}^{n}(y_k - \widehat{y_k})^2}$$

$$MBE = \frac{1}{n}\sum_{k=1}^{n} y_k - \widehat{y_k}$$

where y_k is the actual data point and $\widehat{y_k}$ is the estimated data point.

3.2.3 Benchmark models

ARIMA and deep learning models including MLP, one-dimensional CNN (1DCNN), LSTM, and GRU were implemented and optimised in terms of hyperparameters respectively. A multi-channel 1DCNN (MC-1DCNN) and our proposed DAF-GAN were implemented in the multi-channel settings. Primary model hyperparameters are summarised in Table 1. All the models were trained on a computer equipped with a GeForce RTX 3090.

Table 1. Primary model hyperparameters of models

Method	Primary hyperparameters
ARIMA	p = 24, d = 0, q = 0
MLP	n_fc = 4
1DCNN	n_conv = 3, n_fc = 2
LSTM	n_layers = 2
GRU	n_layers = 2
MC-1DCNN	n_conv = 3, n_fc = 2
DAF-GAN	UNet: n_contract = 6, n_expand = 6 2DCNN: n_contract = 4

Notes: p: Number of Autoregressive Terms; d: Number of Differentiations; q: Number of Moving Average Terms; n_fc: Number of Fully Connected Layers; n_layers: Number of Layers in the Recurrent Unit; n_conv: Number of Convolutional Layers; n_contract: Number of Contracting Blocks; n_expand: Number of Expanding Blocks.

4. RESULTS AND DISCUSSION

Among the compared methods, MLP shows the highest MAPE and RMSE and ARIMA shows the highest CV-RMSE and MAE. The performances of ARIMA and LSTM in terms of MAPE and RMSE are comparable and both are slightly better than MLP. GRU works better than LSTM. While showing the biggest bias, 1DCNN performs the best among the single-channel methods, and its MAPE, CV-RMSE, MAE and RMSE can be reduced by 30.5%, 37.0%, 44.5% and 46.0% respectively.

DAF-GAN shows the least error in terms of all of the metrics. Compared to other methods, its improvements in different metrics fall into the ranges of 21.6%~61.7% (MAPE), 24.2%~65.0% (CV-RMSE), 19.2%~68.3% (MAE), 16.7%~68.1% (RMSE), respectively. DAF-GAN also presents the least forecasting bias by showing the lowest absolute value of MBE compared to other models. Furthermore, it is worth noticing that DAF-GAN shows the least standard deviation across all metrics, which indicates the highest stability across different buildings.

As the global pursuit of energy conservation intensifies and the urgency of precise HVAC energy consumption predictions escalates, the findings of this study stand as a beacon of promise for advancing energy-efficient building automation and management. Furthermore, the potential of this method to propel applied energy research forward and provide insights for decision-making processes underscores its significance in the quest to curtail energy consumption and mitigate environmental impact.

Table 2. Single-building forecasting results (mean ± standard deviation)

Method/Metric	MAPE (%)	CV-RMSE (%)	MAE (kWh)	RMSE (kWh)	MBE (kWh)
ARIMA	23.3 (±7.0)	29.4 (±8.8)	50.5 (±41.3)	63.1 (±51.2)	3.6 (±3.2)
MLP	25.6 (±8.6)	29.0 (±10.4)	49.5 (±41.3)	65.9 (±57.7)	-4.9 (±5.2)
1DCNN	18.0 (±4.3)	21.6 (±5.9)	35.7 (±26.5)	46.9 (±35.5)	7.3 (±8.1)
LSTM	23.1 (±6.8)	28.4 (±8.1)	46.3 (±35.3)	63.2 (±48.1)	1.4 (±7.9)
GRU	21.5 (±8.9)	26.9 (±11.9)	40.8 (±29.2)	55.8 (±41.5)	-3.4 (±17.7)
MC-1DCNN	12.5 (±5.2)	13.6 (±6.3)	19.8 (±9.1)	25.3 (±11.7)	1.0 (±4.2)
DAF-GAN*	**9.8 (±3.3)**	**10.3 (±1.6)**	**16.0 (±8.5)**	**21.0 (±11.0)**	**-0.1 (±1.4)**

5. CONCLUSION

In conclusion, the contribution of this study is two-fold. Firstly, we introduce a novel framework that reconceptualizes forecasting as a problem of conditional generation, facilitating the direct utilization of powerful deep generative models, such as GANs, for energy consumption prediction. Secondly, we present DAF-GAN, a novel multi-channel GAN-based approach for day-ahead building HVAC energy consumption prediction, showcasing relative improvements ranging from 17% to 68% across four distinct error metrics when benchmarked against six traditional and deep learning models. While practical implementation and broader applicability warrant further exploration, our approach holds significant potential for advancing energy-efficient building automation and management. In future research endeavours, emphasis may be placed on the application of DAF-GAN in cross-building forecasting, evaluating performance across various temporal horizons and granularities, and assessing robustness in the face of meteorological uncertainties.

ACKNOWLEDGEMENT

This work was supported by the HKU Foundation Postgraduate Fellowship (awarded to Y.X. Ma) from The University of Hong Kong. The funder has no role in study design, data collection and analysis, interpretation, decision to publish or preparation of the manuscript.

REFERENCE

[1] González-Torres M, Pérez-Lombard L, Coronel J F, et al. A review on buildings energy information: Trends, end-uses, fuels and drivers[J]. Energy Reports, 2022, 8: 626-637.

[2] Box G E P, Jenkins G M, Reinsel G C, et al. Time Series Analysis: Forecasting and Control[M]. John Wiley & Sons, 2015.

[3] Goodfellow I, Pouget-Abadie J, Mirza M, et al. Generative adversarial networks[J]. Communications of the ACM, 2020, 63(11): 139-144.

[4] Fan C, Sun Y, Zhao Y, et al. Deep learning-based feature engineering methods for improved building energy prediction[J]. Applied energy, 2019, 240: 35-45.

[5] Tian C, Li C, Zhang G, et al. Data driven parallel prediction of building energy consumption using generative adversarial nets[J]. Energy and Buildings, 2019, 186: 230-243.

[6] Rumelhart D E, Hinton G E, Williams R J. Learning representations by back-propagating errors[J]. nature, 1986, 323(6088): 533-536.

[7] LeCun Y, Boser B, Denker J, et al. Handwritten digit recognition with a back-propagation network[J]. Advances in neural information processing systems, 1989, 2.

[8] Hochreiter S, Schmidhuber J. Long short-term memory[J]. Neural computation, 1997, 9(8): 1735-1780.

[9] Cho K, Van Merriënboer B, Gulcehre C, et al. Learning phrase representations using RNN encoder-decoder for statistical machine translation[J]. arXiv preprint arXiv:1406.1078, 2014.

[10] Mirza M, Osindero S. Conditional generative adversarial nets[J]. arXiv preprint arXiv:1411.1784, 2014.

[11] Ma X. Human-machine-environment Interaction Investigations Towards Smart Places: Neural Decoding and Building Energy Conservation[J]. HKU Theses Online (HKUTO), 2021.

[12] Wang Z, Oates T. Imaging time-series to improve classification and imputation[J]. arXiv preprint arXiv: 1506.00327, 2015.

[13] Ma Y X, Yu C. Impact of meteorological factors on high-rise office building energy consumption in Hong Kong: From a spatiotemporal perspective[J]. Energy and Buildings, 2020, 228: 110468.

[14] Ma Y X. Short-Term Forecasting of Building Energy Consumption with Deep Generative Learning[C]. Proceedings of the 17th IBPSA Conference. DOI: https://doi.org/10.26868/25222708.2021.31051.

[15] Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation[C]. Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18. Springer International Publishing, 2015: 234-241.

[16] Xiao T, Xu P, He R, Sha H. Status quo and opportunities for building energy prediction in limited data Context—Overview from a competition[J]. Applied Energy, 2022, 305: 117829.