

Generating Contextual Load Profiles Using a Conditional Variational Autoencoder

Wang, Chenguang; Tindemans, Simon H.; Palensky, Peter

DOI

[10.1109/ISGT-Europe54678.2022.9960309](https://doi.org/10.1109/ISGT-Europe54678.2022.9960309)

Publication date

2022

Document Version

Accepted author manuscript

Published in

Proceedings of 2022 IEEE PES Innovative Smart Grid Technologies Conference Europe, ISGT-Europe 2022

Citation (APA)

Wang, C., Tindemans, S. H., & Palensky, P. (2022). Generating Contextual Load Profiles Using a Conditional Variational Autoencoder. In *Proceedings of 2022 IEEE PES Innovative Smart Grid Technologies Conference Europe, ISGT-Europe 2022* (pp. 1-6). (IEEE PES Innovative Smart Grid Technologies Conference Europe; Vol. 2022-October). IEEE. <https://doi.org/10.1109/ISGT-Europe54678.2022.9960309>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Generating Contextual Load Profiles Using a Conditional Variational Autoencoder

Chenguang Wang, Simon H. Tindemans, Peter Palensky
Department of Electrical Sustainable Energy
Delft University of Technology
Delft, The Netherlands
{c.wang-8, s.h.tindemans, p.palensky}@tudelft.nl

Abstract—Generating power system states that have similar distribution and dependency to the historical ones is essential for the tasks of system planning and security assessment, especially when the historical data is insufficient. In this paper, we described a generative model for load profiles of industrial and commercial customers, based on the conditional variational autoencoder (CVAE) neural network architecture, which is challenging due to the highly variable nature of such profiles. Generated contextual load profiles were conditioned on the month of the year and typical power exchange with the grid. Moreover, the quality of generations was both visually and statistically evaluated. The experimental results demonstrate our proposed CVAE model can capture temporal features of historical load profiles and generate ‘realistic’ data with satisfying univariate distributions and multivariate dependencies.

Index Terms—CVAE, generative model, load profiles, synthetic data

I. INTRODUCTION

For power system planning and security assessment, it is of great significance to examine system performance through abundant scenarios [1], [2]. However, when historical data is scarce or larger data sets are required for more precise analysis, it is crucial to building generative models for reproducing unlimited non-repeating data with similar marginal distributions and multivariate dependencies to historical data.

Parametric methods, such as hidden Markov models [3] and *Gaussian mixture models* (GMM), [4] have been utilized to describe historical data patterns. Recently, vine-based copula models have been proposed (e.g., in [5]) to capture marginal distributions and high-dimensional dependencies of historical power system states. However, vine-based copula models are naturally asymmetric and have hard-to-quantify training bias due to sequential model selection.

With the development of machine learning technologies, data-driven generative models, such as *variational autoencoder* (VAE) [6], have been proposed to learn features of high-dimensional historical data and then create ‘unseen’ ones. On this basis, conditional VAE (CVAE) [7] made it possible to generate data under specific conditions. In [8], the impact of the CVAE model’s output noise on its generative performance

has been investigated with a use case of learning and generating snapshots of country-level load states [9]. However, such snapshots of large load aggregations have limited diversity and variability.

In this paper, we bridge the gap by investigating the CVAE model’s capacity to generate synthetic load profiles that are representative of those from a large variety of individual users. Compared to [8], this work aims to generate consumption patterns that are both temporal (instead of spatial) and at a lower aggregation level, where the loads are more stochastic. The main contributions of this paper are as follows:

- We analyze the properties of daily load profiles of an anonymized data set of 5,000 industrial and commercial customers.
- For better training and generation performance, we introduce data split, month condition, and power exchange intensity calculation strategies during data processing.
- We evaluate the performance of the CVAE model under different time (month) and power exchange intensity conditions with both visual and statistical metrics.

II. DATA GENERATION MECHANISM

In this section, a representative multivariate load state generation mechanism was described, based on the *conditional variational autoencoder* (CVAE). The description summarized that in [8].

A. CVAE-based generative model

The CVAE is a neural network architecture that is trained to learn the salient features of historical data by mapping (*encoding*) historical system states onto a lower-dimensional latent space where the latent distribution is approximately normal - and transforming latent vectors back (*decoding*) into a high-dimensional state space [10]. The decoder is used in conjunction with contextual information c to generate representative states (which can be omitted to obtain a regular VAE model). Consequently, the model is able to generate samples with a similar distribution to the historical data by transforming normally distributed samples in the latent space back to the data space. We note that the latent (i.e., hidden) representation of a data point is used solely to facilitate reconstruction and synthesis. It does not need to be imbued with a particular meaning.

This work is supported by the Chinese Scholarship Council.

©2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.




Fig. 1. Schematic of a CVAE model. (a) The structure of a CVAE model for training. (b) The structure of a CVAE model when it's utilized as a generator.

In the training process, the specific structure of the CVAE algorithm is depicted in Fig. 1a. The *encoder* maps the d -dimensional input data x to the code z in the lower-dimensional latent space through k hidden layers H_l^e , $l = 1, \dots, k$. Weight matrices W_l^e , bias vectors b_l^e and the context c are utilized in the encoding process as

$$\begin{aligned} \begin{pmatrix} \mu \\ \sigma \end{pmatrix} &= \begin{pmatrix} W^\mu \\ W^\sigma \end{pmatrix} \left(a(W_k^e(\dots a(W_1^e(x, c) + b_1^e) \dots) + b_k^e) \right) \\ &+ \begin{pmatrix} b^\mu \\ b^\sigma \end{pmatrix}, \end{aligned} \quad (1a)$$

$$z = \mu + \epsilon \odot \sigma, \quad (1b)$$

where a represents an element-wise nonlinear activation function. Vectors μ and σ parameterize an input-dependent normal distribution in the latent space. The output z is sampled accordingly, using ϵ , a vector that is sampled from a standard normal distribution, and the Hadamard product \odot . Mirroring the encoder network, the decoder maps the sampled latent space code z to the d -dimensional data μ' and σ' using

$$\begin{pmatrix} \mu' \\ \sigma' \end{pmatrix} = \begin{pmatrix} W^{\mu'} \\ W^{\sigma'} \end{pmatrix} (\dots a(W_1^d(z, c) + b_1^d) \dots) + \begin{pmatrix} b^{\mu'} \\ b^{\sigma'} \end{pmatrix}, \quad (2)$$

where W_l^d and b_l^d denote weight matrices and bias vectors for decoding, respectively. μ' and σ' parameterize a z -dependent normal distribution in the x space.

After the training process, only the decoder part of the trained CVAE network is utilized to generate data. Latent space codes \tilde{z} are sampled from the standard normal distribution $\mathcal{N}(0, I)$ (see Fig. 1b). Then, data space samples \tilde{x} are sampled from distribution $\mathcal{N}(\mu'(\tilde{z}, c), \sigma'(\tilde{z}, c))$ as $\tilde{x} = \mu' + \epsilon \odot \sigma'$, whose parameters are determined by \tilde{z} and c .

B. Optimization goal

During training, weight matrices W and bias vectors b are updated iteratively to minimize the loss function [10]

$$\mathcal{L} = \mathcal{L}_{D_{KL}} + \mathcal{L}_{Re}. \quad (3)$$

The *Kullback-Leibler loss* $\mathcal{L}_{D_{KL}} = \sum_i D_{KL}(q_\phi(z|x_i)||p(z))$ is the sum over all training data points x_i (assumed i.i.d.)

of the Kullback-Leibler divergence between that point's posterior distribution $q_\phi(z|x_i)$ and the prior distribution $p(z)$ (chosen as the standard normal distribution). The *reconstruction loss* \mathcal{L}_{Re} , representing the negative log-likelihood of reconstructing the inputs x_i via their latent space codes and the decoder that is parameterized by θ , is written as $-\sum_{i=1}^n \mathbb{E}_{Z \sim q_\phi(z|x_i)} [\log P_\theta(x_i|Z)]$. With a constant $\frac{nd}{2} \log 2\pi$ omitted, the \mathcal{L}_{Re} is computed as

$$\mathcal{L}_{Re} \approx \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^d ((x_{i,j} - \mu'_{i,j})^2 / \sigma'^2_{i,j} + \log \sigma'^2_{i,j}), \quad (4)$$

where n denotes the total number of observations used for training. During training, the full-sample sum in loss functions $\mathcal{L}_{D_{KL}}$ and \mathcal{L}_{Re} are replaced by sample batch averages. Inspired by [8], the sample-dependent output noise parameter σ' is co-optimised during training, and the noise $\epsilon \odot \sigma'(\tilde{z}, c)$ is used in the generative process. Also, a weighing factor β was multiplied with $\mathcal{L}_{D_{KL}}$ to adjust the ratio between two losses in (3) as $\mathcal{L} = \beta \mathcal{L}_{D_{KL}} + \mathcal{L}_{Re}$ [11].

III. STUDY DESCRIPTION

We used the CVAE-based generative model described above to generate daily load profiles (24 hours) of individual network connections (i.e., users), conditioned on the month of the year and power the user typically exchanges with the grid. The performance of the model was analyzed using a load data set of 5,000 users. The quality of generations was evaluated visually as a function of conditioning parameters. In addition, performance was validated statistically by measuring univariate distributions and multivariate dependencies. Moreover, an experiment was conducted to test the model capacity of interpolation.

A. Data source

Anonymized historical electricity consumption/generation data of 5,000 industrial and commercial electricity users during 2020 was obtained from Alliander NV [12], a Dutch distribution network owner and operator. The time resolution of the data is 15 minutes. It is worth noting that the data set's time label is *UTC* (Coordinated Universal Time). However, the actual local time for electricity users is *CET* (Central




Fig. 2. Marginal histogram and joint density of historical loads at 10:00 and 21:00 during one year.




Fig. 3. Data processing scheme.

European Time). During standard time and daylight saving time, their time differences are 1 and 2 hours, respectively. The energy data was converted from integer kWh values to average power with multiples of 4 kW. Compared with country-level load profiles [9], the energy consumption of individual users involves more variability and less predictability. Fig. 2 illustrates the large variety of daily profiles, by plotting the marginal histogram and joint density of all historical load profiles at 10:00 and 21:00. Note the logarithmic density used, indicating a large concentration around (relatively) small values. Moreover, data points located in the upper-left and bottom-right corners stand for users can not only consume but also generate energy. All these factors above make it challenging for the CVAE model to capture the load patterns.

B. Data process

The data processing scheme is shown in Fig. 3. The historical data were split, scaled, and conditioned. Three data process strategies used in this study are as follows.




Fig. 4. Training process and its failure. (a) Training process of Kullback-Leibler loss. (b) Training process of reconstruction loss.

1) *Strategy I - Data split*: The historical data were randomly split into training and test sets as blocks of one week with a proportion of 4:1. This strikes a balance between separating individual days (subsequent days are not sufficiently independent) and separating larger blocks (insufficient coverage in the test set).

2) *Strategy II - Month Conditions*: In this study, one of the conditions (contextual information) c is the month of the year. We used the $\sin(\cdot)$ and $\cos(\cdot)$ values of a month as the condition of load data. For a specific month m , its condition was encoded as $\sin(\frac{m}{12} \cdot 2\pi)$ and $\cos(\frac{m}{12} \cdot 2\pi)$. This encoding reflects the continuity and circularity of this feature.

3) *Strategy III - User intensity*: After inspecting historical load profiles, we noticed that some users had relatively regular load profiles, whereas others had irregular behavior with rare consumption or generation spikes. Some connections were only active during a small part of the year. To construct a conditioning feature that represents the ‘size’ of electricity users, we calculated the daily average power exchange by averaging over all non-zero values of the *absolute* power (consumption or generation) for each user and each day. For each customer, this value was averaged over the five days with the largest daily average power exchange to obtain the *user intensity*. The intensity values were used to assign to each customer as a rank order $c \in [0, 1]$. Due to the large range of power values present in the data (see Fig. 2) and the relative scarcity of data with high peak exchange, we trained the model only on profiles of customers with an intensity up to 100 kW. Ultimately, 4,049 users remained, with 1,170,110 and 307,720 load profiles in the training and test sets, respectively. The values were scaled by $1/(100 \text{ kW})$ for training.

C. Training and Data Generation

The parameters of the generative models were tuned for optimal performance. The input and output layers had 96 dimensions (24 hours with 15-minute resolution). Accordingly, 96-dimensional daily load profiles were used for training and generation. The network contained 3 hidden layers in the encoder with dimensions of 800; the bottleneck layer had 12 nodes (12-dimensional latent vector). The decoder also




Fig. 5. 10 randomly sampled historical and generated load profiles and average daily load from customers of various sizes in different months.




Fig. 6. Mean value of historical and generated data in different clusters.

had 3 hidden layers with dimensions of 800. The contextual condition c consists of a 2-dimensional month condition and a 1-dimensional per-user power exchange intensity.

The ReLU activation function was used, except for the generation of μ (μ') and σ (σ') leading up to the bottleneck and output layers. The *adaptive moment estimation* (Adam) weight optimizer [13] was utilized with default settings to iteratively optimize the value of weight matrices W and bias vectors b . The batch size and learning rate parameter α for training were 1,280 and 10^{-5} respectively. The weighting

factor β was set as 8.5. Training and data generation of the model was conducted in Python using `tensorflow` on the Google Colab environment using the GPU option. The training process is shown in Fig. 4. The Kullback-Leibler loss rapidly stabilizes during training. However, the reconstruction loss of the test data set starts to deviate from the training loss and fluctuates strongly after 1,000 training epochs, which indicates an overfitting of training data and general training instability. To find a compromise between loss minimization and generalization capacity of the trained model, 1,000 training epochs

were used in this research. During the generation process, the total, monthly, and per user's amounts of the synthetic data are identical to the training set.

IV. RESULTS

A. Comparison of daily load profiles

To validate the generation capacity of our proposed CVAE model, we first visually inspect the generated contextual load profiles. We define the customers with the first and last 30% of per-user intensities as small and large customers, respectively, and the remaining 40% of users as medium users. In this experiment, we condition the generation of profiles on the months of April and July, and the 'small' and 'large' customer classes (random sampling of c in their respective ranges). Fig. 5 shows the mean value of load profiles under each condition combination (generated versus measured), and 10 randomly sampled load profiles alongside 10 random historical profiles. The mean generated load under each condition combination has a similar curve shape to the training data. Moreover, compared with historical data, the displayed load generations retain randomness and show a sense of realism, indicating that the CVAE model captures temporal features of historical load profiles.

B. Clustering performance

The following experiment compares all historical and generated daily load profiles for a more elaborate test of the distribution of generated load profiles. We first split the training data set into 8 clusters by the K-means algorithm [14], using the squared Euclidean distance metric. Then, we assign the generated and test load profiles to the nearest cluster. The mean values of training, test, and generated loads for each cluster are depicted in Fig. 6a-h, in decreasing order of training data volume. The most voluminous cluster has small average load values. Note that the apparent gap in cluster I is smaller than the resolution of the data. Some clusters correspond to larger loads and generators (mainly solar PV). In all cases, the mean values of profiles assigned to the cluster match well.

C. Marginal distribution comparison

The third experiment compares the cumulative distribution of historical and generated data via different time scales and users of various sizes. The experimental results are shown in Fig. 7; note the discretization of the real measurements, visible in these graphs. Fig. 7a exhibits the cumulative distribution of loads in different months. The CVAE model is able to generate contextual load profiles that follow the monthly distribution variation of historical loads. The hourly comparison of the load depicted in Fig. 7b shows that the curves of generations overlapped with the historical training data, demonstrating quite similar hourly distributions. Moreover, the comparison result shown in Fig. 7c stands for a good capture of load patterns of different customer sizes. Finally, we test the interpolation capacity of the CVAE model. Specifically, we use a virtual month condition (11.5) to generate load profiles, and the result is shown in Fig. 7d. The cumulative distribution




Fig. 7. Cumulative distribution comparison of historical and generated data via different time scales and for users of various sizes.

of the load profiles with a month condition 11.5 lie between the distribution of loads in November and December. This demonstrates that the trained CVAE model can generate data using nonexistent conditions during the training process. Also, these profiles have features of data generated using nearby conditions.

D. Statistical tests

To further test the capacity of the CVAE model to generate realistic load profiles, non-visual statistical tests are implemented to inspect different aspects of generations. Specifically, in this experiment, the Kolmogorov-Smirnov test, autoencoder-based test, and energy test are utilized to examine univariate marginal distributions, point-wise multivariate dependencies, and multivariate dependencies of population, respectively. Interested readers can refer to [8] for more information on these tests. In addition to generations with noise $\epsilon \odot \sigma'(\tilde{z}, c)$ added (these were the data used in previous experiments), we also test the performance of commonly used noise free generations $\mu'(\tilde{z}, c)$ (see also the discussion in [8]).

Evaluating the performance on the Kolmogorov-Smirnov test (Fig. 8a), which assesses the accuracy of the marginal distributions, shows a small difference between the training and test sets, and a similar further difference in the distribution accuracy of the generated data. Comparing the results to those reported in [8] for country-level data, we see a slight degradation of the noisy generator. This could be because the individual load profiles are less smooth than the country-level snapshots, and a relatively large amount of synthetic noise $\epsilon \odot \sigma'(\tilde{z}, c)$ is added to base profiles $\mu'(\tilde{z}, c)$. This can result




Fig. 8. (a) Kolmogorov-Smirnov test. (b) Autoencoder-test. (c) Energy test

in the generation of extreme values, which reduces the test scores.

The autoencoder test trains a separate (regular) autoencoder on the training data. This permits quantification of the quality of individual load profiles. The distributions of reconstruction errors obtained using real and generated data are shown in Fig. 8b. The training and test patterns show similar distributions, and the ‘noisy’ CVAE generates distributions with slightly worse reconstruction errors. However, the ‘noise-free’ variation produces data that is significantly too smooth, resulting in reconstruction errors that are approximately two orders of magnitude lower.

Finally, the energy test quantifies the similarities between high-dimensional *distributions* of profiles. The results in Fig. 8c shows a similar performance between the generated profiles (noisy) and the test data, suggesting good generalization performance. Again, the generated data is a lot more realistic than when no noise is inserted in the output stage (‘noise free’).

V. CONCLUSION AND FUTURE WORK

In this paper, we have investigated the capacity of the CVAE-based model to generate contextual load profiles, randomly selected across a mix of customers: pure loads, pure generators, and mixed load/generators. The load profile generator was trained on data from more than 4,000 industrial and commercial customers, and conditioned on the month of the year and the ‘size’ of the customer, the latter being based on its power exchange with the grid on days with high grid usage. The experimental results demonstrate the model is able to generate visually realistic profiles and perform well on a number of statistical tests. The results also reconfirm the importance of explicitly including (trained) noise in the final stage of the profile generator. In future work, we aim to refine our model to better control the production of extreme load values and more complex dependencies between subsequent moments in time.

ACKNOWLEDGMENT

The authors are grateful to Alliander NV for providing anonymized load data of industrial and commercial users for this study.

REFERENCES

- [1] H. Bloomfield, D. Brayshaw, A. Troccoli, C. Goodess, M. De Felice, L. Dubus, P. Bett, and Y.-M. Saint-Drenan, “Quantifying the sensitivity of European power systems to energy scenarios and climate change projections,” *Renewable Energy*, vol. 164, pp. 1062–1075, 2021.
- [2] P. Panciatichi, G. Bareux, and L. Wehenkel, “Operating in the fog: Security management under uncertainty,” *IEEE Power and Energy Magazine*, vol. 10, no. 5, pp. 40–49, 2012.
- [3] T. Zia, D. Bruckner, and A. Zaidi, “A hidden Markov model based procedure for identifying household electric loads,” in *IECON 2011-37th Annual Conference of the IEEE Industrial Electronics Society*. IEEE, 2011, pp. 3218–3223.
- [4] M.-S. Kang, C.-S. Chen, Y.-L. Ke, C.-H. Lin, and C.-W. Huang, “Load profile synthesis and wind-power-generation prediction for an isolated power system,” *IEEE Transactions on Industry Applications*, vol. 43, no. 6, pp. 1459–1464, 2007.
- [5] I. Konstantelos, M. Sun, S. H. Tindemans, S. Issad, P. Panciatichi, and G. Strbac, “Using vine copulas to generate representative system states for machine learning,” *IEEE Transactions on Power Systems*, vol. 34, no. 1, pp. 225–235, 2019.
- [6] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [7] K. Sohn, H. Lee, and X. Yan, “Learning structured output representation using deep conditional generative models,” *Advances in neural information processing systems*, vol. 28, 2015.
- [8] C. Wang, E. Sharifnia, Z. Gao, S. H. Tindemans, and P. Palensky, “Generating multivariate load states using a conditional variational autoencoder,” *arXiv preprint arXiv:2110.11435*, 2021.
- [9] J. Muehlenpfordt, “Time series,” *Open Power System Data*, 2019. Accessed on: Oct. 3, 2021. [Online]. Available: https://data.open-power-system-data.org/time_series/2019-06-05
- [10] D. P. Kingma and M. Welling, “An Introduction to Variational Autoencoders,” *Foundations and Trends in Machine Learning*, vol. 12, no. 4, pp. 307–392, jun 2019.
- [11] C. P. Burgess, I. Higgins, A. Pal, L. Matthey, N. Watters, G. Desjardins, and A. Lerchner, “Understanding disentangling in β -VAE,” *arXiv preprint arXiv:1804.03599*, 2018.
- [12] nl.wikipedia.org, “Alliander,” *Wikipedia*, 2022. Accessed on: May. 16, 2022. [Online]. Available: <https://nl.wikipedia.org/wiki/Alliander>
- [13] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [14] S. Lloyd, “Least squares quantization in pcm,” *IEEE transactions on information theory*, vol. 28, no. 2, pp. 129–137, 1982.