
EVGen: Adversarial Networks for Learning Electric Vehicle Charging Loads and Hidden Representations

Robert Buechler^{*1} Emmanuel Balogun^{*1} Arun Majumdar¹ Ram Rajagopal²

Abstract

The nexus between transportation, the power grid, and consumer behavior is more pronounced than ever before as the race to decarbonize the transportation sector intensifies. Electrification in the transportation sector has led to technology shifts and rapid deployment of electric vehicles (EVs). The potential increase in stochastic and spatially heterogeneous charging load presents a unique challenge that is not well studied, and will have significant impacts on grid operations, emissions, and system reliability if not managed effectively. Realistic scenario generators can help operators prepare, and machine learning can be leveraged to this end. In this work, we develop generative adversarial networks (GANs) to learn distributions of electric vehicle (EV) charging sessions and disentangled representations. We show that this model structure successfully parameterizes unlabeled temporal and power patterns without supervision and is able to generate synthetic data conditioned on these parameters. We benchmark the generation capability of this model with Gaussian Mixture Models (GMMs), and empirically show that our proposed model framework is better at capturing charging distributions and temporal dynamics.

1. Introduction

The United States, China, and the European Union collectively emit over half of the global greenhouse gas emissions (epa, 2021), highlighting their critical role in combating global climate change. The transportation sector accounts for 29% of total Greenhouse Gas (GHG) emissions in the

U.S, and 14% globally (epa, 2021), making the decarbonization of the transportation sector critical to environmental sustainability. The resurgent urgency to curtail emissions, driven by physical manifestations of climate change today, has motivated governments to create policy that is driving massive Electric Vehicle (EV) adoption. The share of electricity consumption on the power grid attributable to charging EVs is expected to increase up to 40-fold in the United States between 2019 and 2030 (IEA, 2020). EVs accounted for 54% of new car sales in Norway in 2020, up from 42% in 2019 (Reuters, 2021). With the 2035 net-zero carbon goal in the U.S and other electrification targets set by China, UK, and others, we can expect a dense influx of EVs on the roads over the next decade. However, with increased EV penetration, the subsequent massive increase in stochastic power and energy demand from the grid will accelerate the aging of grid assets, complicate optimal operation of power systems, and could require massive upgrades of power systems components. Understanding possible future states of the grid via scenario simulation is paramount to effectively electrifying transportation at scale.

Some work has been done on developing statistical representations of EV charging using Gaussian Mixture Models (GMM) from real charging data (Quirós-Tortós et al., 2018) (Powell et al., 2020). Others have leveraged travel survey data to estimate future EV charging from current driver travel patterns (Negarestani et al., 2016) using traditional statistical approaches. However, few studies utilize deep learning (DL) based approaches, like generative adversarial networks (GAN) (Goodfellow et al., 2014).

GANs were originally developed for single and multi-channel imagery, and GAN architectures continue to be optimized for this application. GANs have also been studied for learning 1D data distributions, including financial timeseries datasets (De Meer Pardo) and medical data anonymization (Bae et al., 2019)(Yoon et al., 2020). Multiple techniques have been developed for conditioning GANs on inputs. A supervised example of this is CGAN (Mirza & Osindero, 2014). InfoGAN (Chen et al., 2016) proposes maximizing the mutual information between latent variables and the output by using a third network, while SCGAN (Li et al., 2019) computes a similarity constraint inside the generator

^{*}Equal contribution ¹Department of Mechanical Engineering, Stanford University, Stanford, California, United States

²Department of Civil and Environmental Engineering, Stanford University, Stanford, California, United States. Correspondence to: Robert Buechler <rbuec@stanford.edu>.

network. As InfoGAN and SCGAN are both unsupervised frameworks, they can be used to learn disentangled representations without labeled data.

2. Dataset

The dataset used in this work is from a real EV charging network at a large tech company campus. The raw data comprises roughly 20000 individual charging events, including variables such as the driver ID, plug-in event ID, charging start and end time, average interval power, energy, driver zip code, and others which are not critical for this paper. Each charging event records power and energy in 15-minute intervals during the charging event. We standardize the length of all data samples into a 24 hour period in 15-minute intervals and assign the charging interval window to a pre-defined standard window, padding the sequence at intervals where each plug-in event has no charging powers. Thus, one data-sample is a 96 X 1 vector representing a charge event. The uniform sequence length simplifies the model and makes it natural to capture representations attributable to charging behavior, battery types, etc. The data is then split between training and test sets with a 95:5 ratio, and is then normalized for training.

3. Experiment

In this work, we utilize a GAN to generate synthetic charging load curves in a specific charging location, conditioned on the learned disentangled representations. We also generate synthetic load curves using a baseline model (GMM) and compare the performance.

3.1. Model Structure

We utilize a GAN variant to learn the probability distribution of the data. Specifically, we use a GAN with Wasserstein loss (WGAN) (Arjovsky et al., 2017) with a gradient penalty to satisfy the Lipschitz constraint (WGAN-GP) (Gulrajani et al., 2017). The gradient penalty term replaces the need for weight clipping, a technique that can limit the models learning ability and lead to undesired behaviour. Wasserstein loss has been shown to improve the stability of training and address mode collapse. We use a combination of stacked convolutional layers and fully connected layers in both the generator (G) and discriminator (D)¹ networks. The specific architecture for each network is shown in Table 1.

We extend the WGAN-GP architecture to learn disentangled representations. Representation learning is important for many tasks and is popularly applied in object detection and

¹This network is also commonly called a critic, since the network is simply trying to maximize the output of real instances, rather than classify between real and fake. We will continue using "discriminator" terminology, although "critic" is also accurate.

image recognition. It is beneficial to be able to perform this task on unlabeled data, since it is more difficult to acquire large open-sourced labeled EV charging datasets. This allows unsupervised algorithms to be trained on anonymized or poorly labeled data. For the unsupervised learning of disentangled representations, we utilize a term (Equation 1) from part of the SCGAN framework. This term encourages the generator to learn the difference between samples of different types and is calculated pair-wise on the training set. This is a simpler method than InfoGAN, since SCGAN shares parameters with G and does not require training a third network. This learning can be done with continuous variables to learn attributes that change smoothly, or discrete variables to denote the boundaries between clearly-defined clusters. The similarity constraint (SC) between condition c and synthetic generated data for continuous conditional variables is written as

$$SC(\mathbf{x}, c) = \frac{1}{N(N-1)} \sum_i \sum_{j \neq i} \left((1 - |c_i - c_j|) sim(\mathbf{x}_i, \mathbf{x}_j) + \frac{|c_i - c_j|}{sim(\mathbf{x}_i, \mathbf{x}_j)} \right) \quad (1)$$

where N is the size of the minibatch, \mathbf{x} is synthetic data from the generator, and $i, j \in [1, N]$. We take $sim(\mathbf{x}_i, \mathbf{x}_j) = \|\mathbf{x}_i - \mathbf{x}_j\|_2$. This term is minimized by the G network and is regularized with λ_{SC} . During training, each c is drawn from $unif(0, 1)$. For discrete variables, the $(1 - |c_i - c_j|)$ term in Eqn 1 is replaced with $\langle \mathbf{c}_i, \mathbf{c}_j \rangle$ and $|c_i - c_j|$ (numerator) is replaced with $(1 - \langle \mathbf{c}_i, \mathbf{c}_j \rangle)$, where \mathbf{c} is a one-hot vector sampled from a multi-nomial distribution. The combination of the SC calculation with WGAN-GP yields a model that we refer to as SC-WGAN-GP. We use the ADAM optimizer, an adaptive learning rate scheduler, and mini-batches with 256 samples per batch.

3.2. Benchmark Model

A Gaussian Mixture Model (GMM) is a convex combination of Gaussians. Each Gaussian represents a group or cluster, and each cluster k , has its own weight α_k , mean μ_k , and variance σ_k^2 , with the mean and variances parametrizing the Gaussian density function for respective clusters. The GMM in its most general form is expressed for a multidimensional variable, where μ_k is a vector and Σ_k is the covariance matrix. The weights α_k , represent probabilities and must sum to one.

$$\sum_{k=1}^K \alpha_k = 1 \quad (2)$$

The parameters for the GMM are obtained via an Expectation Maximization (EM) Algorithm, which seeks to find parameters that maximize the likelihood of the observed data, given the distribution defined by the parameters. If we let our parameters be tuple $\theta = (\alpha, \mu, \Sigma)$. The algorithm

Table 1. Properties of the CNN-based network. N is the size of the minibatch.

Generator		
Layer Type	Parameters	Output Shape
Input (continuous case)	$[z, c] \in \mathbb{R}^{88}$, where $z \sim \text{unif}(0, 1) \in \mathbb{R}^{80}$ and $c \sim \text{unif}(0, 1) \in \mathbb{R}^8$	(N, 88)
Input (discrete case)	$[z, c] \in \mathbb{R}^{88}$, where $z \sim \text{unif}(0, 1) \in \mathbb{R}^{80}$ and $c = e_i \in \mathbb{R}^8$	(N, 88)
Fully Connected	Units=100, LeakyReLU(0.2)	(N, 150)
Expand Dimension	—	(N, 1, 150)
Conv1D	Filters = 32, kernel size= 5, padding = 'replicate', LeakyReLU(0.2)	(N, 32, 150)
Conv1D	Filters = 16, kernel size= 5, padding = 'replicate', LeakyReLU(0.2)	(N, 16, 150)
Conv1D	Filters = 8, kernel size= 5, padding = 'replicate', LeakyReLU(0.2)	(N, 8, 150)
Conv1D	Filters = 1, kernel size= 5, padding = 'replicate', LeakyReLU(0.2)	(N, 1, 150)
Squeeze Dimension	—	(N, 150)
Fully Connected	Units=125, LeakyReLU(0.2)	(N, 125)
Fully Connected	Units=100, LeakyReLU(0.2)	(N, 100)
Fully Connected	Units=96, LeakyReLU(0.2)	(N, 96)
Discriminator / Critic		
Layer Type	Parameters	Output Shape
Input	—	(N, 96)
Expand Dimension	—	(N, 1, 96)
Conv1D	Filters = 32, kernel size= 5, padding = 'replicate', LeakyReLU(0.2)	(N, 32, 96)
MaxPool	PoolSize = 2	(N, 32, 48)
Conv1D	Filters = 16, kernel size= 5, padding = 'replicate', LeakyReLU(0.2)	(N, 16, 48)
MaxPool	PoolSize = 2	(N, 16, 24)
Conv1D	Filters = 8, kernel size= 5, padding = 'replicate', LeakyReLU(0.2)	(N, 8, 24)
Flatten	—	(N, 192)
Fully Connected	Units=50, LeakyReLU(0.2)	(N, 50)
Fully Connected	Units=15, LeakyReLU(0.2)	(N, 15)
Fully Connected	Units=1, LeakyReLU(0.2)	(N, 1)

solves the optimization problem to maximize the likelihood,

$$\mathbf{Q}(\theta^*, \theta) = \sum_{N=1}^N \sum_{k=1}^K \gamma(z_{nk}) [\ln \alpha_k + \ln \mathcal{N}(x_n | \mu_k, \Sigma_k)] - \lambda \left(\sum_{k=1}^K \alpha_k - 1 \right) \quad (3)$$

given the observed data true distribution. γ represents the probability of a particular cluster k , given the observed data. The last term includes the Lagrange multiplier λ , which serves to constrain the sum of the Gaussian weights to one. To develop the GMM, we postulate that a simplified EV charging session is completely described by charging session start time, charging session duration, and charging session average load intensity. Thus, the GMM models the prior and the joint probability given each prior. This means it learns the probability for each cluster and captures the joint probability of charge start times, charging duration, and load intensities of charge sessions, given some cluster k with weight (prior probability) α_k . GMMs use the popular Expectation Maximization (EM) algorithm to iteratively solve for a minimum.

3.3. Results

We train our benchmark (the GMM) and SC-WGAN-GP models and compare performance. For this experiment, we choose a generator with an 80-dimensional input latent

space z from $\text{unif}(0, 1)$. For representation learning, we train one model with eight continuous variables and another model with eight discrete variables. The SC-WGAN-GP trains for around 4000 epochs and takes about 9hrs to converge on a NVIDIA K80 GPU, although it can produce good (visually, qualitative) fake samples after around 500 epochs. The training curves for both networks can be viewed [here](#) for the continuous SC-WGAN-GP and [here](#) for the discrete SC-WGAN-GP. The D network, SC, and GP terms are stable and converge.

For the GMM, we selected the cluster count to be 1000, as this was the main parameter we tuned for this application. The maximum number of iterations for the Expectation Maximization (EM) algorithm was set to 50000 and the convergence tolerance was 10^{-6} . The GMM displays improving performance with the number of clusters, up to a limit, at which it starts to overfit. We find the optimal cluster count N to be around the chosen value of 1000. The GMM converges in about 6 - 10 minutes, which is much faster than the SC-WGAN-GP. Performance statistics for various numbers of clusters are shown in Figure A1 in the supplementary materials.

3.4. Discussion

It is imperative that the generated data can display mode diversity (the ability of the model to capture multiple modes

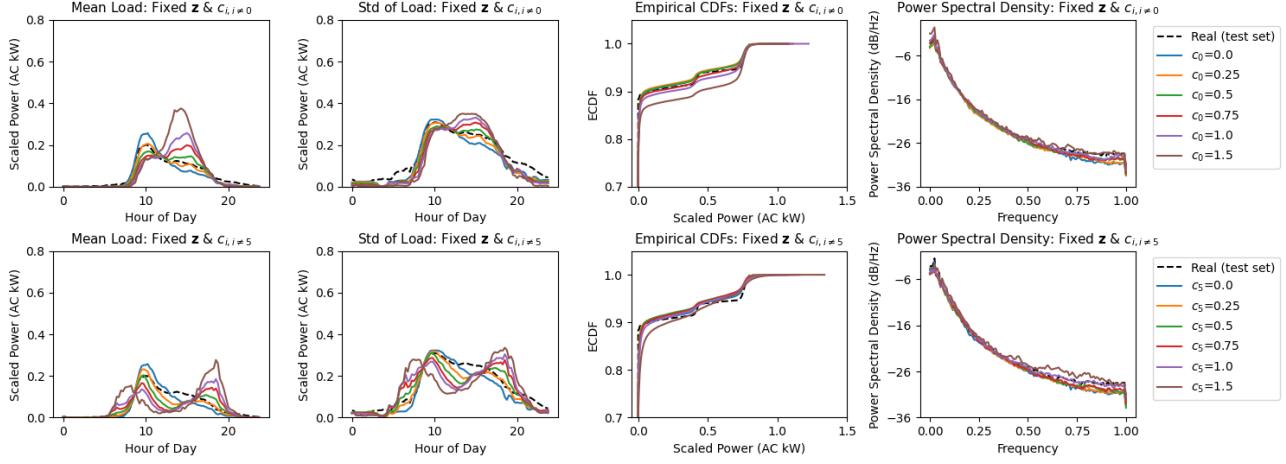


Figure 1. Statistics of SC-WGAN-GP (continuous) generator output. Only two continuous variables c_0 (top) and c_5 (bottom) are shown. Each line represents a different continuous variable value.

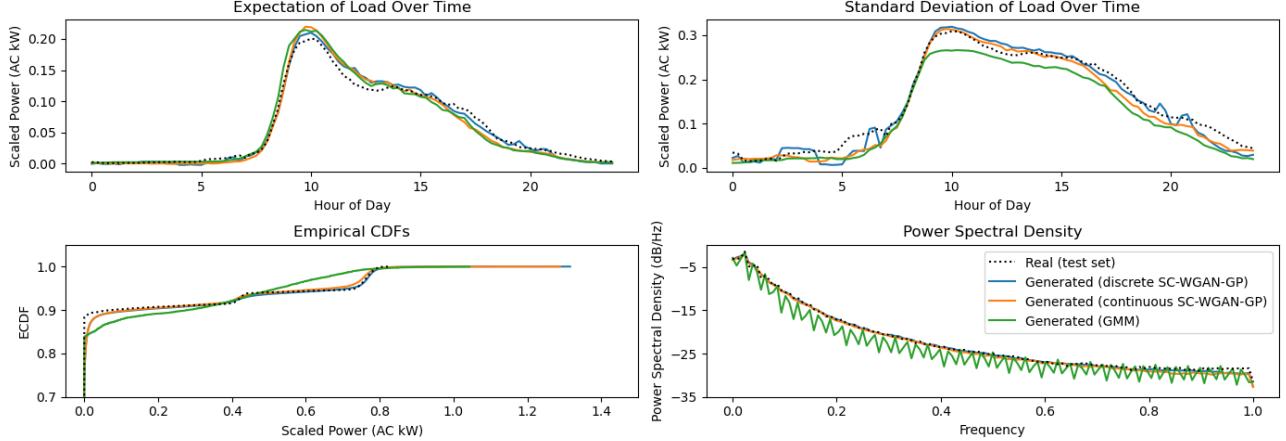


Figure 2. Summary statistics for the three networks: continuous SC-WGAN-GP, discrete SC-WGAN-GP, and GMM. The GAN architectures most notably outperform the GMM when considering the the CDF and PSD of the signals.

from the training set, not just one) and statistical resemblance (the generated output should have similar statistical properties as the real distribution). It is insignificant to compare single samples from each distribution; the entire distributions must be analyzed as a whole. To analyze the probability of certain-valued loads being present in each distribution, the cumulative distribution function (CDF) is derived empirically from the real and generated sample loads. To analyze the periodic component and the temporal fluctuations present in each distribution, the power spectral density is calculated, as is shown along with the CDF in Figure 2.

For the GAN, we observe a trade-off between generator performance and representation learning. If λ_{SC} is high,

the generator may not converge, and if λ_{SC} is low, the generator may not learn to capture representations. λ_{SC} therefore requires careful tuning, and may benefit from dynamic scheduling; although we saw no immediate successes of such a structure in our experiments, it should be explored further. The authors of SCGAN implement the pair-wise loss calculation as a nested iteration loop for each minibatch parameter update step, and cite their method as being 3-5x slower than InfoGAN. We implement a fully vectorized version of the SC pair-wise calculation that is at least 50x faster than SCGAN.

Figure 1 shows two of the eight continuous variables learned by SC-WGAN-GP. c_0 appears to increase the power level of charging sessions in the afternoon in exchange of de-

creasing sessions in the morning with higher values of c_0 . c_5 appears to shift consumption out to later in the evening while drastically reducing midday consumption at higher c_5 values. We found that other variables (not shown) learned various other temporal and power-level patterns. The value for 1.5 for each variable demonstrates model extrapolation since values were sampled from $unif(0, 1)$ during training. This is significant for this application because it offers the capability to test future scenarios where certain variables are expected to change, on large unlabeled datasets.

Figure A2 in the supplementary materials show the discrete representations learned by SC-WGAN-GP. It is clear that the discrete variables are grouping load shapes by time and charging duration.

An inherent limitation of the GMM model is its inability to capture transient charging dynamics for each generated charging session, since the model assumes an average steady load for each session. One possible mitigation for this will be to sample from one or very few similar GMM clusters multiple times, and then take the average for each time-interval. With this idea, the GMM still fails to produce realistic loads but instead, produces stepped load shapes. Because the GMM model is fast and performs okay, it is better suited for long-term planning of power infrastructure, but not real-time operation and simulation. For operations and control, transient dynamics are important as the power system and all components that connect to it are predominantly dynamic. Additionally, because the GMM models steady loads, information on the frequency driven properties that might aid the learning of subtle and latent representative groups like the EV car type, model, or battery chemistry, are lost.

4. Conclusion

In this work, we built a framework that successfully captures hidden representations in an EV charging dataset in an unsupervised manner. The results show that our proposed framework, SC-WGAN-GP, is able to generate statistically realistic EV load curves, and can learn hidden temporal and power patterns in the training set better than more traditional modeling methods. Also importantly, the latent variables are intuitive and have physical meaning; for instance, some variables group drivers with vehicles that charge fast, while some other variables capture temporal charging behavior from the distribution. This shows promising signs of developing representative EV charging scenarios for regions globally, which will be instrumental for grid infrastructural planning/operation. DL-based methods like GANs could be used for EV load generation for demand response or vehicle-to-grid applications, where accurate short-term (second to minute scale) representations of load behaviour is critical. We propose the following future improvements to the model

architecture and training procedure:

- *Constrained SC loss.* Some continuous variables learned very similar patterns during training, resulting in redundant learning. We aim to constrain SC loss to learn dissimilar embeddings to minimize statistical redundancy in variables.
- *More interpretable model extrapolation.* The SC-WGAN-GP (continuous) is able to extrapolate learned representations past the distribution seen in the training set. However, the form of this extrapolation is unknown, and may not be linear. Mapping these learned representations to a more interpretable space would help in real applications, i.e. increasing the value of a variable that has learned battery size.
- *Regional Learning.* Although the data used in this work is limited to a single region, we can train our model to learn distributions for different regions by conditioning the GAN outputs on locational variables.
- *Address bias in training data.* As our training data was limited to a single company campus location, we aim to test the model’s generalizability by training with public EV charging station data.

Acknowledgements

Thank you to Siobhan Powell and Gustavo Vianna Cezar for their advice and support, including suggestions related to benchmark models and data sources. Thank you to our reviewers for providing helpful insight.

References

Global greenhouse gas emissions data, Mar 2021. URL <https://www.epa.gov/ghgemissions/global-greenhouse-gas-emissions-data>.

Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein gan. *arXiv:1701.07875 [cs, stat]*, Dec 2017. URL <http://arxiv.org/abs/1701.07875>. arXiv: 1701.07875.

Bae, H., Jung, D., and Yoon, S. Anomigan: Generative adversarial networks for anonymizing private medical data. *arXiv:1901.11313 [cs]*, Jan 2019. URL <http://arxiv.org/abs/1901.11313>. arXiv: 1901.11313.

Chen, X., Duan, Y., Houthooft, R., Schulman, J., Sutskever, I., and Abbeel, P. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *arXiv:1606.03657 [cs, stat]*, Jun 2016. URL <http://arxiv.org/abs/1606.03657>. arXiv: 1606.03657.

De Meer Pardo, F. Enriching financial datasets with generative adversarial networks — tu delft repositories. URL <https://repository.tudelft.nl/islandora/object/uuid%3A51d69925-fb7b-4e82-9ba6-f8295f96705c>.

Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial networks. *arXiv:1406.2661 [cs, stat]*, Jun 2014. URL <http://arxiv.org/abs/1406.2661>. arXiv: 1406.2661.

Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. Improved training of wasserstein gans. *arXiv:1704.00028 [cs, stat]*, Dec 2017. URL <http://arxiv.org/abs/1704.00028>. arXiv: 1704.00028.

IEA. *Global EV Outlook 2020*. 2020. URL <https://www.iea.org/reports/global-ev-outlook-2020>.

Li, X., Chen, L., Wang, L., Wu, P., and Tong, W. Scgan: Disentangled representation learning by adding similarity constraint on generative adversarial nets. *IEEE Access*, 7: 147928–147938, 2019. ISSN 2169-3536. doi: 10.1109/ACCESS.2018.2872695.

Mirza, M. and Osindero, S. Conditional generative adversarial nets. *arXiv:1411.1784 [cs, stat]*, Nov 2014. URL <http://arxiv.org/abs/1411.1784>. arXiv: 1411.1784.

Negarestani, S., Fotuhi-Firuzabad, M., Rastegar, M., and Rajabi-Ghahnavieh, A. Optimal sizing of storage system in a fast charging station for plug-in hybrid electric vehicles. *IEEE Transactions on Transportation Electrification*, 2(4):443–453, 2016. doi: 10.1109/TTE.2016.2559165.

Powell, S., Kara, E. C., Sevlian, R., Cezar, G. V., Kilicotte, S., and Rajagopal, R. Controlled workplace charging of electric vehicles: The impact of rate schedules on transformer aging. *Applied Energy*, 276:115352, 2020. ISSN 0306-2619. doi: <https://doi.org/10.1016/j.apenergy.2020.115352>. URL <https://www.sciencedirect.com/science/article/pii/S0306261920308643>.

Quirós-Tortós, J., Espinosa, A. N., Ochoa, L. F., and Butler, T. Statistical representation of ev charging: Real data analysis and applications. In *2018 Power Systems Computation Conference (PSCC)*, pp. 1–7, 2018. doi: 10.23919/PSCC.2018.8442988.

Reuters. Norway's ev sales rise to record 54 URL <https://europe.autonews.com/automakers/norways-ev-sales-rise-record-54-market-share>.

Yoon, J., Drumright, L. N., and van der Schaar, M. Anonymization through data synthesis using generative adversarial networks (ads-gan). *IEEE Journal of Biomedical and Health Informatics*, 24(8):2378–2388, Aug 2020. ISSN 2168-2194, 2168-2208. doi: 10.1109/JBHI.2020.2980262.

Supplementary Materials

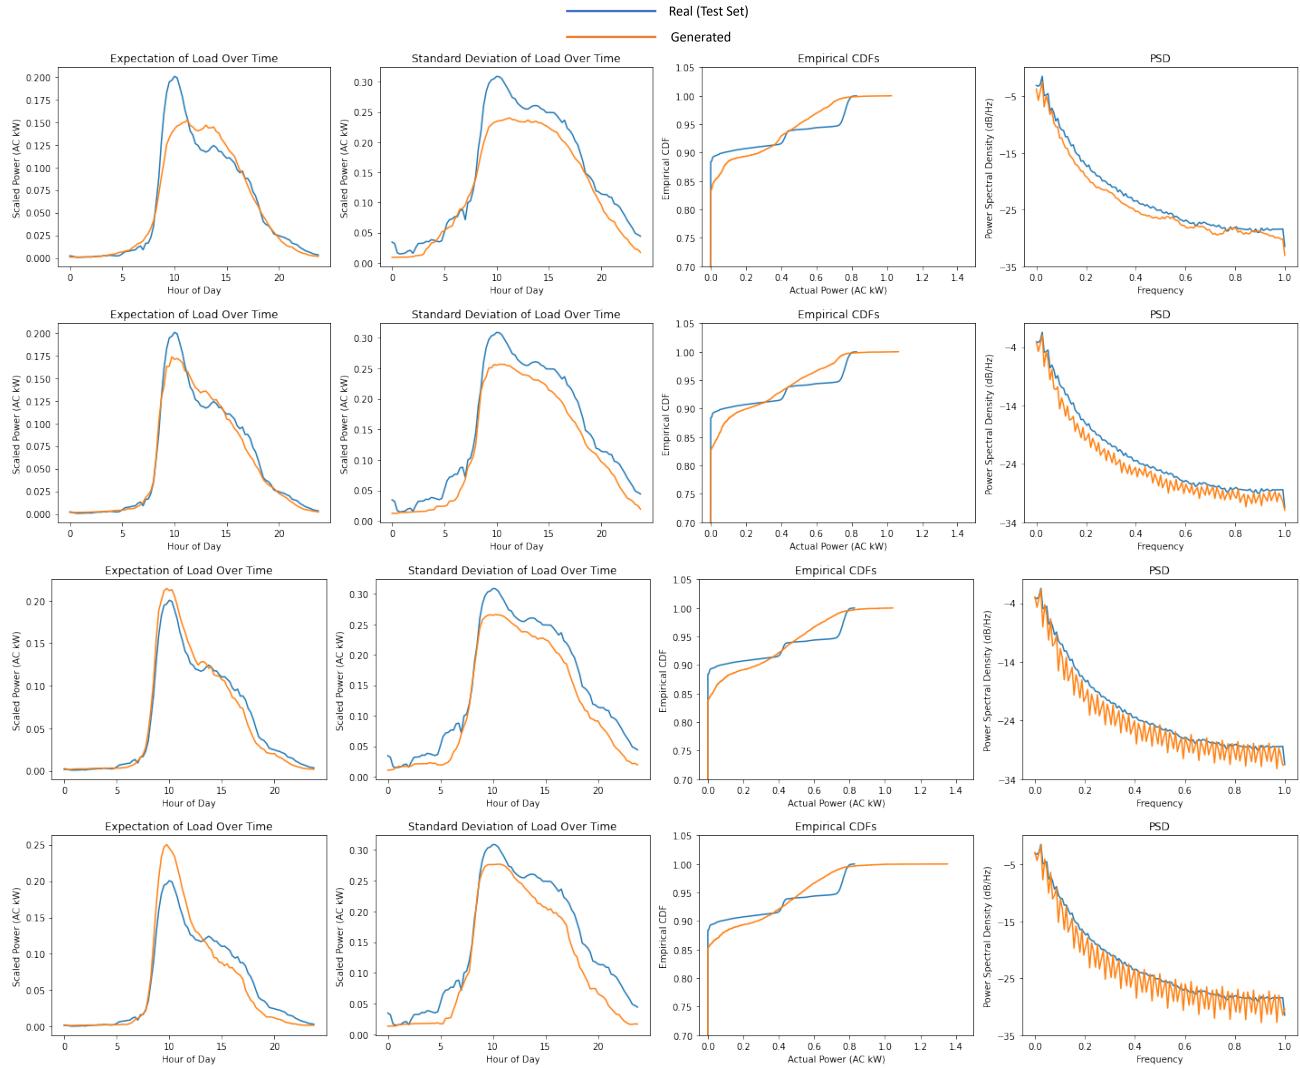


Figure A1. Summary statistics for GMM Number of clusters (N = 125, 500, 1000, 2000. Top to bottom). The GMM performance degrades for N = 2000.

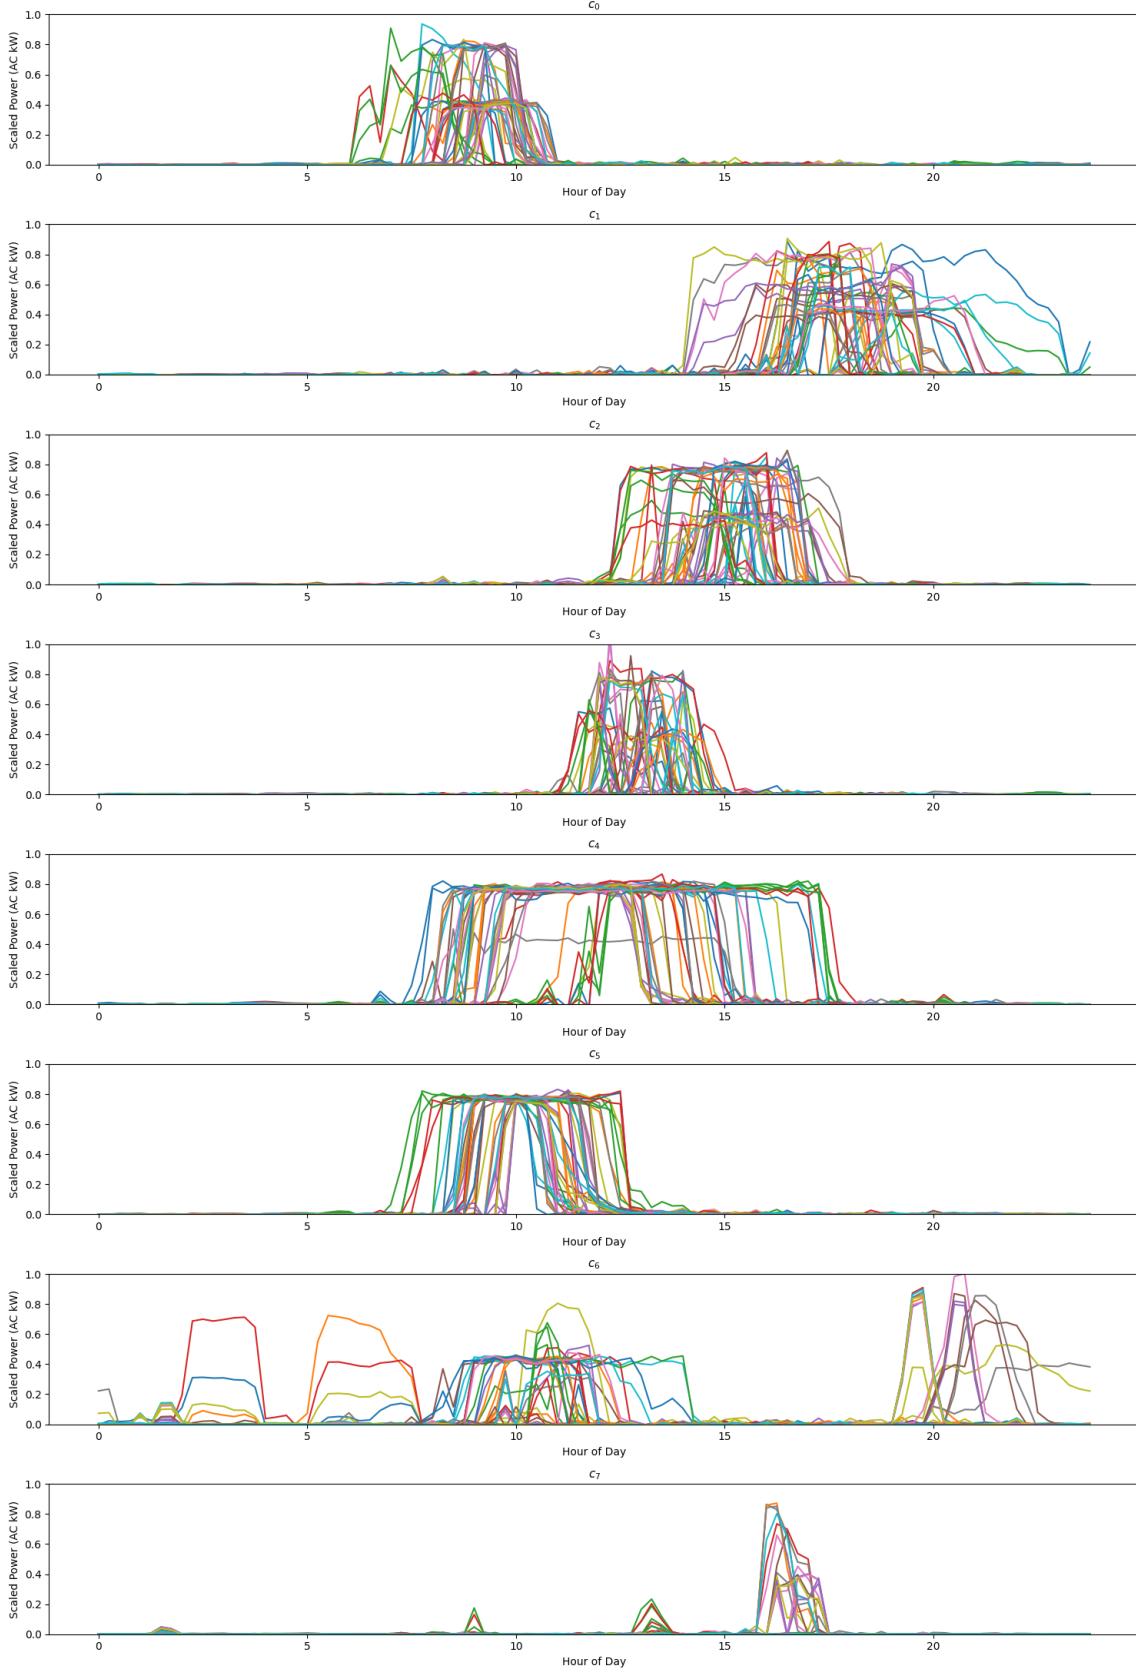


Figure A2. Sample synthetic data generated from the discrete SC-WGAN-GP (50 samples per plot). All plots were generated using the same latent space z . Each row includes samples from a different discrete variable.