
Challenges in Gaussian Processes for Non Intrusive Load Monitoring

Aadesh Desai Gautam Vashishtha Zeel B Patel Nipun Batra

IIT Gandhinagar, India

{desai.aadesh,gautam.pv,patel.zeel,nipun.batra}@iitgn.ac.in

Abstract

Non-intrusive load monitoring (NILM) or energy disaggregation aims to break down total household energy consumption into constituent appliances. Prior work has shown that providing an energy breakdown can help people save up to 15% of energy. In recent years, deep neural networks (deep NNs) have made remarkable progress in the domain of NILM. In this paper, we demonstrate the performance of Gaussian Processes (GPs) for NILM. We choose GPs due to three main reasons: i) GPs inherently model uncertainty; ii) equivalence between infinite NNs and GPs; iii) by appropriately designing the kernel we can incorporate domain expertise. We explore and present the challenges of applying our GP approaches to NILM.

1 Introduction

Prior work has shown that providing an energy breakdown: per-appliance energy consumption can empower users to reduce energy consumption by up to 15% [1]. The most accurate way to get an energy breakdown is via “disaggregating” the household total energy consumption measured via a smart meter, which is a single point of instrumentation per home. This line of work is often called energy disaggregation or non-intrusive load monitoring (called non-intrusive, as there is no instrumentation or intrusion in the home) [2].

The initial work in the domain involved heuristics and edge-detection-based methods [3]. In the past decade, graphical model variants such as the factorial hidden Markov model were proposed [4]. However, in the recent past, deep neural networks have outperformed conventional methods. These models leverage the advances in NNs across various domains such as attention.

In this work, we explore Gaussian Processes (GPs) for energy disaggregation. Our motivation is threefold: i) GPs inherently model uncertainty. Recent work on applying Bayesian neural networks for NILM suggests poor calibration performance [5]; ii) given the equivalence between infinite NNs and GPs [6], we may expect to get comparable or better performance for GPs; iii) GPs can incorporate domain expertise via kernel design [7].

2 Methods

Given the mains (aggregate) power consumption at time \bar{y}_t at time t , our aim is to estimate the power \hat{y}_t^n for the n^{th} appliance. We now introduce our GP variants. In this paper, we fit a separate GP per appliance. We summarise our models in Table 1. We also note that due to prohibitive training and testing time and memory requirement for exact GPs, we use sparse GPs [8] for all our GP variants. We also use automatic relevance determination (A.R.D.) for models that take a vector input.

Point-to-Point: In our first model, we fit a GP with the aggregate power at time t as the input and the appliance power draw \hat{y}_t^n as the output.

Sequence-to-Point: Inspired by the state-of-the-art NN model for NILM [9], in our next model, the input to the GP is a window of main readings around time t , while the output from the GP is the

Model	Input features	Kernel	Output
Point-to-point GP	\bar{y}_t	Matern (5/2)	\hat{y}_t^n
Sequence-to-point GP	Sequence of $2k+1$ mains readings ($\bar{y}_{t-k:t+k}$)	Matern (5/2) with A.R.D.	\hat{y}_t^n
Feature Based GP	$\bar{y}_t, \max(\bar{y}_{t-k:t+k}), \min(\bar{y}_{t-k:t+k}),$ $\text{mean}(\bar{y}_{t-k:t+k}), \text{Kurtosis}(\bar{y}_{t-k:t+k}),$ Difference: $\bar{y}_t - \bar{y}_{t-1}, \text{range}(\bar{y}_{t-k:t+k})$	Matern (5/2) with A.R.D.	\hat{y}_t^n

Table 1: Kernels, the input features and output for each GP model

appliance reading at time t . With such a model, we hope to capture the context of the aggregate reading. As an example, aggregate power reading of 100 Watts followed by multiple readings of 0 Watts can indicate an appliance turning off, whereas the aggregate power reading of 100 Watts followed by similar power readings indicate the appliances remaining in the same state.

Feature Based: Our next model is inspired from the time series based models which incorporate advanced statistical features [10]. We create a latent representation of the $2k+1$ length sequence of points that were previously used in the sequence to point model using the following features shown in Table 1. Such statistical features may provide hints to our model. For example, a high minimum value from the $2k+1$ length sequence may indicate that multiple appliances are On during that window (higher base load). Similarly, a low value of range (difference between maximum and minimum power) may indicate that the appliance has not changed state.

Linear Kernel: The central problem we face in all of our models is the incorrect prediction of the OFF state (we get an positive offset in all cases). In order to create a trainable parameter that can counter this offset, we incorporate the “range” feature inside a Linear Kernel in order to provide explicit hints to our model about its minimum and maximum state values. We train sequence-to-point as well feature-based model with this modification.

3 Evaluation

3.1 Dataset

We use the publicly available REDD [11, 12] dataset for our evaluation. The REDD dataset consists of data from six homes over several weeks. In this work we use data from three homes similar to prior work [12]. The mains data is at one second frequency, and the appliance data is at three seconds frequency. We resample both the mains and appliance data to a minute resolution, as it is a common practice in the domain [13]. The mains data consists of apparent power (vector sum of active and reactive power), whereas the appliance data consists of active power.

3.2 Experimental Setup

In this paper, we create an artificial aggregate instead of a true aggregate, similar to prior work [13]. The artificial aggregate is created by summing up the Refrigerator, Dishwasher and Microwave power at a given time. We also note that using the artificial aggregate simplifies the problem.

Additionally, to test the robustness of our models to changes in the power distribution during test time, we modify the test aggregate power by adding a constant 100 Watts load. We leave the true test appliance power as it is. This experiment helps simulate the scenario when an unseen load is present in the test home.

For the sequence-to-point model, we have chosen the sequence length as 99 as used by prior work [14]. Our GP kernel parameters are initialized from a standard normal distribution. We perform hyperparameter tuning using grid search to find optimal inducing points, learning rate and epochs for each model. The metric used to evaluate is mean absolute error (MAE) = $\frac{1}{n} \sum_{i=1}^n |\hat{y}_i - y_i|$. We also report the uncertainty using the following metrics: i) CE: is calculated as the absolute difference between x and the number of samples falling within the $x\%$ confidence interval. We report CE for a 95% confidence interval (Hatalis et al. 2017); ii) MSLL: Mean Standardized Log Loss (MSLL) is

Model	Dataset					
	Artificial Mains			Artificial Mains + 100		
	MAE	MSLL	CE(95%)	MAE	MSLL	CE(95%)
Point-to-point	18.43	9.41	0.036	78.89	38.68	0.033
Sequence-to-point	8.76	5.26	0.049	92.18	47.91	0.007
Sequence-to-point + Linear Kernel	10.46	5.50	0.047	88.62	43.44	0.013
Statistical Features	7.60	5.13	0.045	84.24	41.15	0.005
Statistical Features + Linear Kernel	7.13	5.03	0.046	73.48	33.52	0.030

Table 2: Metrics for different GP methods.

an average of log pdf over all the test points considering the predictive distribution (Rasmussen and Williams 2005). Table 2 shows the main results of our experiments. In all our approaches we use leave-one-home-out cross-validation [12].

In the interest of space and time, we only present the results on refrigerator in this paper. Our work is fully reproducible and can be found at https://github.com/Aadesh-1404/nilm_gp.

3.3 Results and Analysis

Our main result is shown in Table 2. We first discuss the results on artificial aggregate followed by artificial aggregate + test bias dataset. Overall, we find our statistical features model with the linear kernel on the range performs the best.

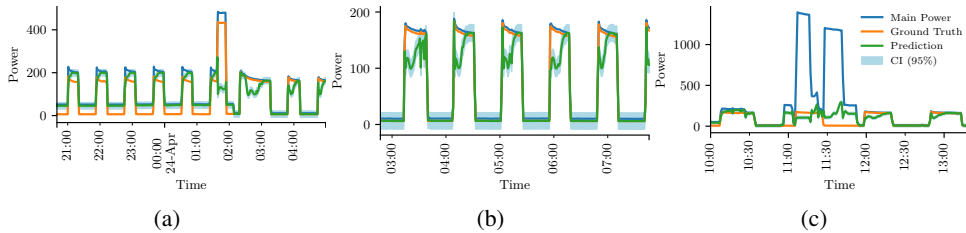


Figure 1: Predictions from the Point-to-point model

3.3.1 Artificial aggregate

Point-to-Point: Figure 1 shows the predictions for our point-to-point GP model for three different windows of time. In Figure 1(a), we can note that the ground truth fridge power is 0 Watts when the fridge is OFF (e.g. 22:30 to 23:00 hours). However, our model predicts the power as 40 Watts. This can be explained using 2(a), which shows the training appliance vs. aggregate data and the GP fit. 3 patterns can be observed in this figure: 1) the $y=x$ line denotes the case when the only refrigerator is turned ON, 2) $y=K$ (around 200 Watts) line denotes the case when multiple devices are turned ON along with the refrigerator and 3) $y=0$ line denotes when the refrigerator power is 0, but other appliances are ON. Thus, we can explain our model’s non-zero (approximately 40 Watts) prediction using these 3 states. The constant fluctuation between these three states drives the prediction of the model to an intermediate value in this range predicting to 0 Watts. Figure 1 (b) depicts the case when the refrigerator is the only active appliance (and thus, mains power is equal to refrigerator power). We observe a certain amount of dip at the rising edge of every cycle. We can explain this from the scatter plot in Figure 2 (c), which illustrates the inadequate number of data points in the rising edge, faltering the model’s predictive abilities in that region.

At first, reaching the ON state (188 W), the value falls to a lower value corresponding to mains (188 W) as depicted in Figure 2 (a), then again rises to the value corresponding to mains (150 W). Figure 1 (c) illustrates the prediction of the model at higher values of power. It can be observed that there is vast fluctuation in the predictions in this region which can be reasoned by the density plots shown in Figure 2 (b). The test readings have a higher density around 200 Watts compared to the train readings and thus the GP model is unable to make accurate predictions in this power range.

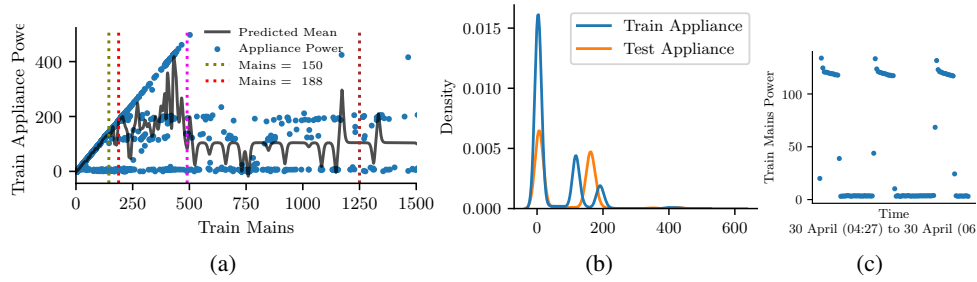


Figure 2: Point-to-point Analysis: (a) Prediction of the GP with the training data (b) Density plot for appliance power distribution (c) scatter plot of train mains power

Sequence-to-Point: This model performs significantly better than the point-to-point approach. One of the earlier mentioned challenges of the point-to-point model was predicting the ON state values correctly when the fridge was the only ON appliance. Figure 3 (b) shows that the Sequence-to-point model performs significantly better than the point-to-point model (Figure ??). Also, from 3 (c), we can note that when observing sudden/unseen changes in the training mains, the GP prediction are smooth, which was not the case in Point-to-point Figure 1 (b). However, as seen from 3 (a), though predictions for defrost stage improve, the model still fails to accurately predict the OFF stage.

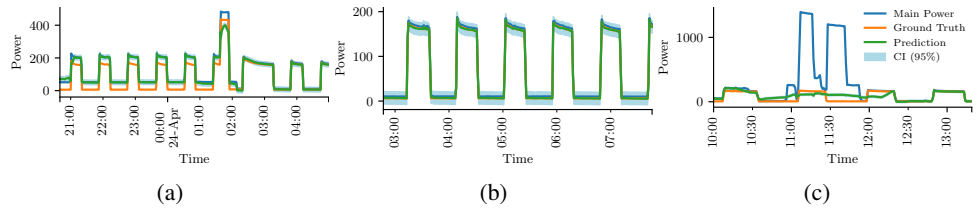


Figure 3: Predictions for the Sequence-to-sequence model

Features to Point: We can observe from 4(b) that this model performs significantly better than point-to-point GP and also beats sequence-to-point GP. 4(c) also illustrates that we are able to predict the defrost state better using the model. From, 4(a), we conclude Range along with difference as the important features from the ARD plot. However, we can also see that this model still fails in correctly predicting the OFF state 4(c).

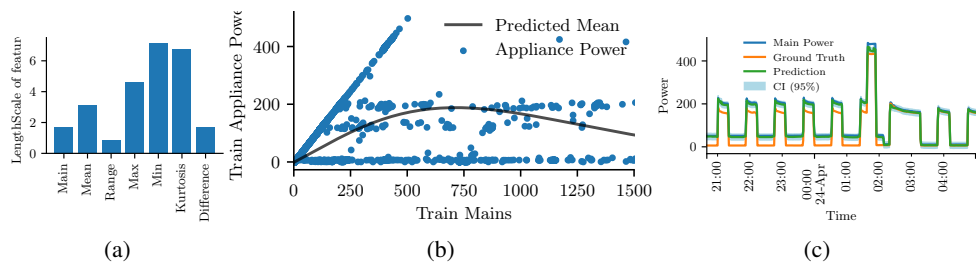


Figure 4: Feature based GP: (a) A.R.D plot, (b) Mean Prediction (c) Predictions showing failure

3.3.2 Artificial aggregate + Test Bias

Point-to-point - Fig 5 (a) illustrated on adding a bias, the results fail significantly, as the mains is shifted it simply gives predictions for per point shift as per the predicted mean in Fig 5 (a).

Sequence-to-point - this model performs worse than Point-to-point. This can be explained by the covariance matrix being formed from 99-time steps which all add a uniform bias, making the model make incorrect predictions with more certainty and thus the inferior nature of the prediction. To mitigate the bias, we create a learnable parameter by adding the range (max-min) of the mains feature and incorporating it in a linear kernel. This strategy improved predictions.

Features to point: This model fails as expected, but the results are better than Sequence-to-point, this is because the bias is not uniform across all the features and the distributed bias gives better predictions. On using range with linear kernel it is able to compensate for the bias (Fig 5(d)).

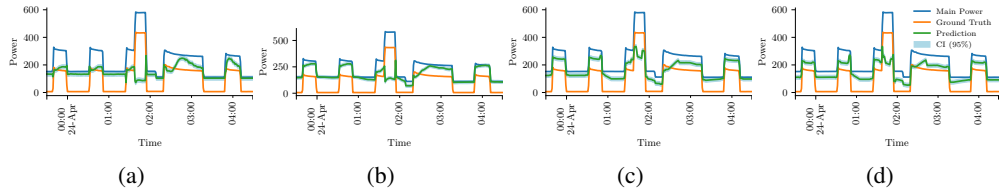


Figure 5: GP Output on adding bias: (a) Point-to-point (b) Sequence-to-point (c) Feature based (d) Feature based + Linear Kernel

We quantify uncertainty in prediction by plotting the reliability curves (in the Appendix) and using the expected calibration error [5]. The reliability curve (calibration curve) can help visually understand the model’s calibration while the expected calibration error (ECE) quantitatively measures the model calibration. A lower ECE value indicates a better calibration. We have computed ECE values for our models in Table 3.

Model	Dataset	
	Artificial Mains	Artificial Mains + 100
Point-to-point	0.196	0.483
Sequence-to-point	0.229	0.496
Sequence-to-point + Range (Linear Kernel)	0.085	0.493
Statistical Features	0.269	0.493
Statistical Features + Range (Linear Kernel)	0.279	0.486

Table 3: Expected Calibration Error

4 Conclusion & Future Work

In this work, we investigate GPs for energy NILM. Energy disaggregation is a complex multi-trend dataset which is difficult to predict with point-to-point approaches. The sequence-to-point model performed better due to access to a larger bin of information but lacked in terms of obscure latent representation of the feature space. Explicit featurization in the feature-to-point model performed the best due to access to concise information and a semantically rich latent space. We are optimistic that significant progress lies ahead in this area of study. Advanced featurization and explicit kernel designing can be useful for appropriately disaggregating power for all the appliances. Further, we can explore Deep GP models, Deep Kernel Learning and MultiOutput GPs for NILM.

References

- [1] Sarah Darby et al. The effectiveness of feedback on energy consumption. *A Review for DEFRA of the Literature on Metering, Billing and direct Displays*, 486(2006):26, 2006.
- [2] D Mariano-Hernández, L Hernández-Callejo, A Zorita-Lamadrid, O Duque-Pérez, and F Santos García. A review of strategies for building energy management system: Model predictive control, demand side management, optimization, and fault detect & diagnosis. *Journal of Building Engineering*, 33:101692, 2021.
- [3] G.W. Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870–1891, 1992.
- [4] Jack Kelly and William Knottenbelt. Neural nilm: Deep neural networks applied to energy disaggregation. In *Proceedings of the 2nd ACM international conference on embedded systems for energy-efficient built environments*, pages 55–64, 2015.

- [5] Vibhuti Bansal, Rohit Khoiwal, Hetvi Shastri, Haikoo Khandor, and Nipun Batra. “i do not know”: Quantifying uncertainty in neural network based approaches for non-intrusive load monitoring. In *Proceedings of the 9th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, 2022.
- [6] Radford M Neal. Priors for infinite networks. In *Bayesian Learning for Neural Networks*, pages 29–53. Springer, 1996.
- [7] Zeel B Patel, Palak Purohit, Harsh M Patel, Shivam Sahni, and Nipun Batra. Accurate and scalable gaussian processes for fine-grained air quality inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 12080–12088, 2022.
- [8] Michalis Titsias. Variational learning of inducing variables in sparse gaussian processes. In *Artificial intelligence and statistics*, pages 567–574. PMLR, 2009.
- [9] Chaoyun Zhang, Mingjun Zhong, Zongzuo Wang, Nigel Goddard, and Charles Sutton. Sequence-to-point learning with neural networks for non-intrusive load monitoring. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- [10] Michael W Coughlin, Kevin Burdge, Dmitry A Duev, Michael L Katz, Jan Van Roestel, Andrew Drake, Matthew J Graham, Lynne Hillenbrand, Ashish A Mahabal, Frank J Masci, et al. The ztf source classification project–ii. periodicity and variability processing metrics. *Monthly Notices of the Royal Astronomical Society*, 505(2):2954–2965, 2021.
- [11] J Zico Kolter and Matthew J Johnson. Redd: A public data set for energy disaggregation research. In *Workshop on data mining applications in sustainability (SIGKDD)*, San Diego, CA, volume 25, pages 59–62. Citeseer, 2011.
- [12] Hetvi Shastri and Nipun Batra. Neural network approaches and dataset parser for nilm toolkit. In *Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pages 172–175, 2021.
- [13] Nipun Batra, Jack Kelly, Oliver Parson, Haimonti Dutta, William Knottenbelt, Alex Rogers, Amarjeet Singh, and Mani Srivastava. Nilmtk: An open source toolkit for non-intrusive load monitoring. In *Proceedings of the 5th international conference on Future energy systems*, pages 265–276, 2014.
- [14] Nipun Batra, Rithwik Kukunuri, Ayush Pandey, Raktim Malakar, Rajat Kumar, Odysseas Krystalakos, Mingjun Zhong, Paulo Meira, and Oliver Parson. Towards reproducible state-of-the-art energy disaggregation. In *Proceedings of the 6th ACM international conference on systems for energy-efficient buildings, cities, and transportation*, pages 193–202, 2019.

A Appendix

A.1 Calibration

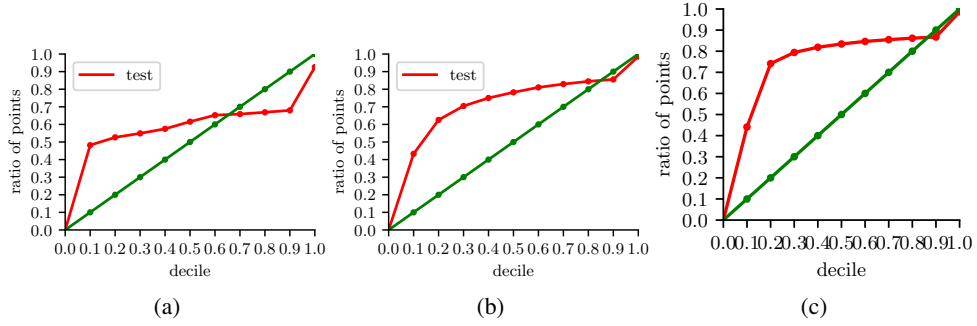


Figure 6: Calibration Plots: (a) Point-to-point GP, (b) Sequence-to-point GP (c) Feature Based GP

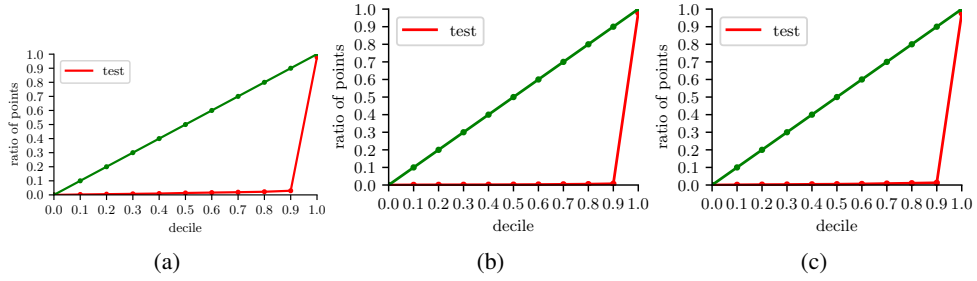


Figure 7: Calibration Plots Added Bias: (a) Point-to-point GP, (b) Sequence-to-point GP (c) Feature Based GP

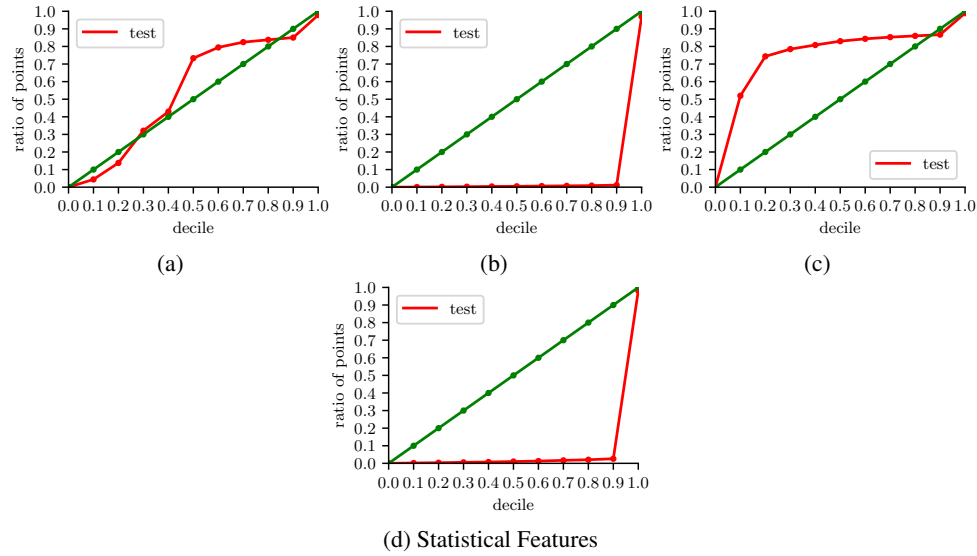


Figure 8: Calibration Plots Linear: (a) Seq-to-point + Linear GP, (b) Bias: Seq-to-point + Linear GP, (c) Feature Based + Linear GP, (d) Bias: Feature Based + Linear GP