
Sliced-Wasserstein-based Anomaly Detection and Open Dataset for Localized Critical Peak Rebates

Julien Pallage*

Dept. of Electrical Engineering
Polytechnique Montréal, GERAD & Mila

Bertrand Scherrer

Data Science, Hilo
Hydro-Québec

Salma Naccache

Data Science, Hilo
Hydro-Québec

Christophe Bélanger

Open Data | Data and Analytics
Hydro-Québec

Antoine Lesage-Landry

Dept. of Electrical Engineering
Polytechnique Montréal, GERAD & Mila

Abstract

In this work, we present a new unsupervised anomaly (outlier) detection (AD) method using the sliced-Wasserstein metric. This filtering technique is conceptually interesting for MLOps pipelines deploying machine learning models in critical sectors, e.g., energy, as it offers a conservative data selection. Additionally, we open the first dataset showcasing localized critical peak rebate demand response in a northern climate. We demonstrate the capabilities of our method on synthetic datasets as well as standard AD datasets and use it in the making of a first benchmark for our open-source localized critical peak rebate dataset.

1 Introduction

Québec, Canada, stands as an outlier in the electrical grid decarbonization paradigm. As there is a global tendency to invest in intermittent renewable energy sources to decarbonize grids worldwide, Québec, while not totally stranger to this trend, is generally mostly carbon-free thanks to its impressive hydroelectric power capacity. One of its main challenges comes from its northern climate, its reliance on electric baseboard heating systems for residential heating, and its unrestrictive home insulation policies [1]. During glacial winter days, as electric heaters are all running at once, and peak consumption hours hit, Québec’s hydropower capacity can be reached [2]. To accommodate winter peak power needs, Hydro-Québec, the state-owned company in charge of electric power generation, transmission, and distribution, must operate its only on-grid thermal power plant and import electricity from neighbouring provinces and states. These imports are usually expensive and produce much more greenhouse gas emissions in comparison to local energy sources [3].

To remediate this issue without solely relying on the deployment of new generation and transmission infrastructures, demand response (DR) initiatives have flourished in the province [4]. Demand response can be defined as changes in normal electrical power consumption patterns by end-use customers in response to a signal, e.g., financial incentives or control setpoints [5]. With Québec’s long tradition of fixed electricity rates, one of its main DR mechanisms is critical peak rebates (CPR). Residential customers enrolled in a CPR program receive financial compensation during pre-specified time periods, referred to as challenges, for reducing their energy consumption with respect to their expected baseload [6]. CPR programs are purely voluntary and virtually penalty-free. They thus depend on consumers’ goodwill, motivation, and sensitivity. Yet, as we have seen in our work, CPR events can be powerful tools for shifting power consumption before and after each event.

*Corresponding author: julien.pallage@polymtl.ca.

We work with a variation of CPR, viz., localized CPR (LCPR), in which the events are called for localized relief in the grid instead of being cast for the whole system. LCPR can diversify the types of services offered by typical CPR programs, e.g., they can alleviate stress on local equipment like substation transformers [7] or they can be paired with generation forecasts of distributed energy resources (DERs), e.g., roof solar panels and private windfarms, to balance demand and generation during peak hours. LCPR is highly underexplored in the literature and is a valuable application to benchmark trustworthy machine learning (ML) models. Indeed, the higher spatial granularity, the critical aspect of the task, the dependence on behavioural tendencies, and the lower margin for error require the deployment of forecasting models that offer performance guarantees [8], robustness to noise [9], interpretability [10], physical constraints satisfaction [11], a sense of prediction confidence [12, 13], or a combination of them [14]. Being able to predict and utilize localized peak-shaving potential in the electrical grid, through programs similar to LCPR, could accelerate the integration of DERs and, specifically for Québec, phase out its dependence on fossil energy-based imports. To the best of our knowledge, no published open-source datasets are showcasing either LCPR or CPR schemes, in a northern climate.

As with most smart grid applications, dataset quality can highly influence the performance of ML models when used for training. The adversarial properties stemming from the amalgam of numerical errors, noise in sensor readings, telemetry issues, meter outages, and unusual extreme events can disrupt the prediction quality of ML models. Anomaly detection (AD) and outlier filtering are thus primordial in a reliable ML pipeline [15]. Unsupervised AD methods are preferable as they do not need human-made labels and their hyperparameters can be tuned simultaneously with other ML models' included in the loop. Popular unsupervised AD methods [16] include local outlier factor (LOF) [17], isolation forest [18], k-nearest neighbours (KNN) [19], connectivity-based outlier factor [20], and one-class support vector machine (SVM) [21]. These methods either use clustering or local density to assign an outlier score. We are interested in optimal transport-based (OT) metrics for AD. Reference [22] proposed a Wasserstein generative adversarial network for AD tasks and authors from [23] have designed a differentiable training loss function using the Wasserstein metric for deep classifiers. To the best of our knowledge, no unsupervised OT method has been proposed yet for AD.

In this work, we address the lack of open-source datasets showcasing CPR demand response mechanisms in northern climates by releasing two years of aggregated consumption data for customers participating in an LCPR program in Montréal, Québec, Canada². These customers are spread in three adjacent distribution substations. With it, we hope to stimulate research on trustworthy ML models for demand response applications. We also address the challenges of training ML models with unfiltered smart grid data by proposing a new simple unsupervised outlier filter leveraging the Sliced-Wasserstein distance [24]. We showcase the performance of this filter on standard anomaly detection datasets and leverage it to present a benchmark for the prediction of the localized energy consumption of LCPR participants on our open-source dataset.

2 Localized critical peak rebates in Québec

Québec's CPR and LCPR programs have been carried out under the banner of Hilo, a division of Hydro-Québec in charge of DR aggregation. Hilo, calls a DR *challenge* a day ahead of the event, and users choose their degree of participation for the next day. Hilo subscribers are equipped with smart thermostats and their respective heating setpoints are controlled by Hilo, according to an agreed-upon strategy, throughout the event. With the Hilo mobile application and different connected objects, customers can program the response of their house to different *scenes*. For example, they can choose the heating setpoint of each thermostat when there is a DR event, or when nobody is home. When notified, smart thermostats pre-heat the house to augment user comfort during events which typically last 4 hours: either between 6 AM and 10 AM or 5 PM and 9 PM. A maximum of 30 CPR events can be called per year. LCPR is an additional program currently under testing. Testers can be asked for up to 10 extra LCPR events. Rewards are proportional to the total amount of energy shaved during the event, i.e., heating-related curtailment and others, with respect to their estimated baselines [25].

²Available at <https://github.com/jupall/lcpr-data>, as well as, <https://donnees.hydroquebec.com/explore/dataset/consommation-clients-evenements-pointe>

3 Open dataset

Description The dataset we share contains the aggregated hourly consumption of 197 anonymous LCPR testers located in three substations. Additional hourly weather data and LCPR information are also present. Table 1 details the features and the label. Note that we also provide cyclical encoding of temporal features, e.g., month, day of the week, and hour. We remark that outliers and anomalies are present in the dataset because of metering and telemetry issues or even blackouts, e.g., an aberrant (and impossible) 32.2 MWh energy consumption is registered at some point. We refer readers to Appendix A for additional analyses and visualizations of the dataset’s features and labels.

Open data initiatives at Hydro-Québec The yearly energy demand in Québec should double by 2050, requiring an additional production of 60 TWh by 2035 and the addition of 8 to 9 GW of peak power to Québec’s 38 GW capacity [26]. Québec’s energy landscape must evolve rapidly to meet the increasing demand and DR is a powerful tool to mitigate infrastructure growth.

Hydro-Québec recognizes the importance of sharing data to foster research, innovation, and informed decision-making. Through its platform³, it provides historical and real-time data on energy production and consumption in Québec as well as other datasets, e.g., geolocated hydrometeorological data from its remote weather stations. This new dataset is Hydro-Québec’s first dedicated to DR research.

By supporting data democratization, Hydro-Québec encourages the emergence of a culture of transparency and openness in the energy sector which is needed to accelerate the energy transition. This approach aims to stimulate researchers, citizens, and businesses in their desire to participate in the design of innovative applications for a more sustainable, interoperable, reliable, and safe power grid.

4 Sliced-Wasserstein Filter

The Wasserstein distance is a metric that provides a sense of distance between two distributions. Also called earth mover’s distance, it can be conceived as the minimal effort it would take to displace a pile of a weighted resource to shape a specific second pile [27]. The order- t Wasserstein distance with respect to some norm $\|\cdot\|$ between distributions \mathbb{U} and \mathbb{V} is defined as:

$$W_{\|\cdot\|,t}(\mathbb{U}, \mathbb{V}) = \left(\min_{\pi \in \mathcal{J}(\mathbb{U}, \mathbb{V})} \int_{\mathcal{Z} \times \mathcal{Z}} \|\mathbf{z}_1 - \mathbf{z}_2\|^t d\pi(\mathbf{z}_1, \mathbf{z}_2) \right)^{1/t},$$

where $\mathcal{J}(\mathbb{U}, \mathbb{V})$ represents all the possible joint distributions π between \mathbb{U} and \mathbb{V} , \mathbf{z}_1 and \mathbf{z}_2 are the marginals of \mathbb{U} and \mathbb{V} , respectively, and \mathcal{Z} is the set of all possible values of \mathbf{z}_1 and \mathbf{z}_2 [9]. In general, computing the Wasserstein distance is of high computational complexity [28] except for some special cases. For example, the Wasserstein distance for one-dimensional distributions has a closed-form solution that can be efficiently approximated. The sliced-Wasserstein (SW) distance is a metric that makes use of this property by calculating infinitely many linear projections of the high-dimensional distribution onto one-dimensional distributions and then computing the average of the Wasserstein distance between these one-dimensional representations [24, 28]. Interestingly, it possesses similar theoretical properties to the original Wasserstein distance while being computationally tractable [29]. The order- t SW distance’s approximation under a Monte Carlo scheme is defined as:

$$SW_{\|\cdot\|,t}(\mathbb{U}, \mathbb{V}) \approx \left(\frac{1}{L} \sum_{l=1}^L W_{\|\cdot\|,t}(\mathcal{R}\mathbb{U}(\cdot, \theta_l), \mathcal{R}\mathbb{V}(\cdot, \theta_l))^t \right)^{1/t},$$

where $\mathcal{R}\mathbb{D}(\cdot, \theta_l)$ is a single projection of the Radon transform of distribution function \mathbb{D} over a fixed sample $\theta_l \in \mathbb{S}^{d-1} = \{\boldsymbol{\theta} \in \mathbb{R}^d | \theta_1^2 + \theta_2^2 + \dots + \theta_d^2 = 1\} \forall l \in \llbracket L \rrbracket$. Interested readers are referred to references [24, 28] for in-depth mathematical explanations.

We now utilize this approximation of the sliced-Wasserstein distance to formulate a simple outlier filter. Consider an empirical distribution $\hat{\mathbb{P}}_N = \frac{1}{N} \sum_{i=1}^N \delta_{\mathbf{z}_i}(\mathbf{z})$, where $\delta_{\mathbf{z}_i}(\cdot)$ is a Dirac delta function assigning a probability mass of one at each known sample $\mathbf{z}_i \in \mathbb{R}^d$. Consider the notation $\hat{\mathbb{P}}_{N-1}^{-\mathbf{z}_i}$ to denote a variation of $\hat{\mathbb{P}}_N$ in which we remove sample \mathbf{z}_i . As the weight of each sample is identical and because the SW distance compares distributions of equal weights, we propose an outlier filter by

³[https://www.hydroquebec.com/documents-data/](https://www.hydroquebec.com/documents-data/open-data/)

using a simple voting system in which we compare the SW distance between the empirical distribution minus the outlier candidate and the empirical distribution minus a random sample. Let $O \subseteq \mathcal{D}$ and $O^c \subseteq \mathcal{D}$ denote the sets of outlier and inlier samples, respectively, under the perspective of the filter and let $\mathcal{D} = O \cup O^c \subset \mathbb{R}^d$ be the available dataset. A vote is positive if the distance between the two empirical distributions exceeds the threshold $\epsilon > 0$. A sample is labelled an outlier if the proportion of positive votes is greater than the threshold p .

$$O = \left\{ \mathbf{z}_i \in \mathcal{D} \mid p \leq \frac{1}{n} \sum_{\mathbf{z}_j \sim \hat{\mathbb{P}}_{N-1}^{\mathbf{z}_i}} \mathbb{1} \left(SW_{\|\cdot\|, t}(\hat{\mathbb{P}}_{N-1}^{\mathbf{z}_i}, \hat{\mathbb{P}}_{N-1}^{\mathbf{z}_j}) \geq \epsilon \right), \forall i \in \llbracket N \rrbracket \right\}, \quad (\text{SWAD})$$

where $\mathbb{1}(\cdot)$ denotes an indicator function, n is the number of points used for the vote, $p \in [0, 1]$ is the voting threshold required to label a sample as an outlier, and $\llbracket N \rrbracket \equiv \{1, 2, \dots, N\}$.

This filter is interesting for several reasons. It is unsupervised, purely analytical, and uses a well-known explainable mathematical distance to filter data points that seem out-of-sample under the SW metric. The intuition is that we remove samples that are costly in the transportation plan when compared to other random samples of the same distribution. We can use the voting percentage to measure the algorithm’s confidence in its labelling. It is also parallelizable, as seen in our implementation provided on our GitHub page⁴. Numerical results and figures are presented in Appendix B.

This method does not scale with large datasets as is: the SW distance computational burden increases as bigger distributions are compared. We propose a smart splitting method to accelerate the procedure, but we remark that the transportation plan between a distribution minus a sample and the same distribution minus another sample can be roughly approximated by the Euclidian distance between the two removed samples. As such, we introduce a fast Euclidian approximation for $W_{\|\cdot\|, 1}$:

$$O \approx \left\{ \mathbf{z}_i \in \mathcal{D} \mid p \leq \frac{1}{n} \sum_{\mathbf{z}_j \sim \hat{\mathbb{P}}_{N-1}^{\mathbf{z}_i}} \mathbb{1} (\|\mathbf{z}_i - \mathbf{z}_j\|_2 \geq \eta), \forall i \in \llbracket N \rrbracket \right\}, \quad (\text{FEAD})$$

where η is the threshold of the Euclidian distance. This method is not as accurate as the first proposal to filter out-of-sample data points. But, because (SWAD) and (FEAD) use the same principle, they share similar classification patterns when ϵ and η are tuned accordingly, as can be seen in Figure 7.

5 Benchmark

We now propose a first benchmark on our LCPR dataset. Our goal is to predict the aggregated hourly consumption at each substation in winter when peak demand is critical. To follow the literature [30], and propose a simple yet meaningful benchmark, we implement a Gaussian process [31] in a rolling horizon fashion. Samples dating from before 2023-12-15 are used for hyperparameter tuning while those between 2023-12-15 and 2024-04-15, during the most recent winter DR season, are used for testing. We train one model per week and the training window size is a hyperparameter to be tuned. The SW filter is used to preprocess data in the tuning phase and increase generalization in testing. Our method is interesting because it can filter clear out-of-sample points while avoiding sparse LCPR events that other methods could consider as local outliers. We refer interested readers to our GitHub page for more details. Testing mean average errors (MAE) and root mean squared errors (RMSE) are presented in Table 2 for each substation. See Appendix C for additional figures.

6 Closing remark

In this work, we present a new unsupervised outlier filter by leveraging the sliced-Wasserstein metric. This filter is interesting for MLOps integration on applications where global outliers may be adversarial to the prediction quality of trained models, e.g., smart grid data. We also hope to stimulate research on trustworthy ML models in critical sectors by releasing the first open dataset showcasing localized critical peak rebate demand response schemes in a northern climate. This dataset has a strong potential for benchmarking of such models as it opens a window to a real-world critical application where accuracy and robustness are equally important. To get the ball rolling, we provide a first benchmark by tuning simultaneously our SW filter and a Gaussian process.

⁴<https://github.com/jupall/swfilter>

References

- [1] T. Gerbet, “Performance énergétique des bâtiments : Québec manque de courage, déplorent des experts,” *Radio-Canada*, Sep 2023.
- [2] J. Whitmore and P.-O. Pineau, “État de l’énergie au Québec,” 2024, préparé pour le gouvernement du Québec.
- [3] Hydro-Québec, *Une énergie renouvelable pour un Québec durable, Rapport sur le développement durable*, 2023.
- [4] F. Pelletier and A. Faruqi, “Does dynamic pricing work in a winter-peaking climate? A case study of Hydro Quebec,” *The Electricity Journal*, vol. 35, no. 2, p. 107080, 2022.
- [5] M. H. Albadi and E. F. El-Saadany, “Demand Response in Electricity Markets: An Overview,” in *2007 IEEE Power Engineering Society General Meeting*, 2007, pp. 1–5.
- [6] A. Mercado, R. Mitchell, S. Earni, R. Diamond, and E. Alschuler, “Enabling interoperability through a common language for building performance data,” *Proceedings of the 2014 ACEEE Summer Study on Energy Efficiency in Buildings*, 2014.
- [7] F. Li, I. Kocar, and A. Lesage-Landry, “A Rapid Method for Impact Analysis of Grid-Edge Technologies on Power Distribution Networks,” *IEEE Transactions on Power Systems*, vol. 39, no. 1, pp. 1530–1542, 2024.
- [8] A. Venzke and S. Chatzivasileiadis, “Verification of neural network behaviour: Formal guarantees for power system applications,” *IEEE Transactions on Smart Grid*, vol. 12, no. 1, pp. 383–397, 2021.
- [9] R. Chen and I. C. Paschalidis, “Distributionally robust learning,” *Foundations and Trends® in Optimization*, vol. 4, no. 1-2, pp. 1–243, 2020.
- [10] L. H. Gilpin, D. Bau, B. Z. Yuan, A. Bajwa, M. A. Specter, and L. Kagal, “Explaining Explanations: An Approach to Evaluating Interpretability of Machine Learning,” *CoRR*, vol. abs/1806.00069, 2018.
- [11] G. S. Misyris, A. Venzke, and S. Chatzivasileiadis, “Physics-informed neural networks for power systems,” in *2020 IEEE Power & Energy Society General Meeting (PESGM)*, 2020, pp. 1–5.
- [12] L. V. Jospin, H. Laga, F. Boussaid, W. Buntine, and M. Bennamoun, “Hands-on bayesian neural networks—a tutorial for deep learning users,” *IEEE Computational Intelligence Magazine*, vol. 17, no. 2, pp. 29–48, 2022.
- [13] A. G. Wilson and P. Izmailov, “Bayesian deep learning and a probabilistic perspective of generalization,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 4697–4708, 2020.
- [14] J. Pallage and A. Lesage-Landry, “Wasserstein Distributionally Robust Shallow Convex Neural Networks,” *arXiv preprint arXiv:2407.16800*, 2024.
- [15] J. Stiasny, S. Chevalier, R. Nellikkath, B. Sævarsson, and S. Chatzivasileiadis, “Closing the loop: A framework for trustworthy machine learning in power systems,” *Proceedings of 11th Bulk Power Systems Dynamics and Control Symposium (IREP 2022)*, 2022.
- [16] M. Goldstein and S. Uchida, “A Comparative Evaluation of Unsupervised Anomaly Detection Algorithms for Multivariate Data,” *PLOS ONE*, vol. 11, no. 4, pp. 1–31, 04 2016.
- [17] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, “LOF: identifying density-based local outliers,” *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, p. 93–104, 2000.
- [18] F. T. Liu, K. M. Ting, and Z.-H. Zhou, “Isolation Forest,” in *2008 Eighth IEEE International Conference on Data Mining*, 2008, pp. 413–422.

- [19] F. Angiulli and C. Pizzuti, “Fast Outlier Detection in High Dimensional Spaces,” *Proceedings of the Sixth European Conference on the Principles of Data Mining and Knowledge Discovery*, vol. 2431, pp. 15–26, 08 2002.
- [20] J. Tang, Z. Chen, A. W.-C. Fu, and D. W.-L. Cheung, “Enhancing Effectiveness of Outlier Detections for Low Density Patterns,” in *Proceedings of the 6th Pacific-Asia Conference on Advances in Knowledge Discovery and Data Mining*, ser. PAKDD ’02. Berlin, Heidelberg: Springer-Verlag, 2002, p. 535–548.
- [21] M. Amer, M. Goldstein, and S. Abdennadher, “Enhancing one-class support vector machines for unsupervised anomaly detection,” *Proceedings of the ACM SIGKDD Workshop on Outlier Detection and Description*, p. 8–15, 2013.
- [22] M. Ducoffe, I. Haloui, and J. S. Gupta, “Anomaly detection on time series with Wasserstein GAN applied to PHM,” *International Journal of Prognostics and Health Management*, vol. 10, no. 4, 2019.
- [23] Y. Wang, W. Sun, J. Jin, Z. Kong, and X. Yue, “WOOD: Wasserstein-Based Out-of-Distribution Detection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 2, pp. 944–956, 2024.
- [24] N. Bonneel, J. Rabin, G. Peyré, and H. Pfister, “Sliced and Radon Wasserstein barycenters of measures,” *Journal of Mathematical Imaging and Vision*, vol. 51, pp. 22–45, 2015.
- [25] Hilo. (2023, 09) Residential. [Online]. Available: <https://www.hiloenergie.com/en-ca/>
- [26] Hydro-Québec, *Action Plan 2035, Towards a Decarbonized and Prosperous Québec*, 2023.
- [27] V. M. Panaretos and Y. Zemel, “Statistical aspects of wasserstein distances,” *Annual review of statistics and its application*, vol. 6, no. 1, pp. 405–431, 2019.
- [28] S. Kolouri, K. Nadjahi, U. Simsekli, R. Badeau, and G. Rohde, “Generalized sliced wasserstein distances,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [29] N. Bonnotte, “Unidimensional and evolution methods for optimal transportation,” Ph.D. dissertation, Université Paris Sud-Paris XI; Scuola normale superiore (Pise, Italie), 2013.
- [30] Y. Weng and R. Rajagopal, “Probabilistic baseline estimation via Gaussian process,” in *2015 IEEE Power & Energy Society General Meeting*, 2015, pp. 1–5.
- [31] C. Williams and C. Rasmussen, “Gaussian processes for regression,” *Advances in neural information processing systems*, vol. 8, 1995.
- [32] C. Spearman, “The Proof and Measurement of Association Between Two Things,” *American Journal of Psychology*, vol. 15, pp. 88–103, 1904.
- [33] T. Chen and C. Guestrin, “XGBoost: A Scalable Tree Boosting System,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD ’16. New York, NY, USA: Association for Computing Machinery, 2016, p. 785–794.
- [34] S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS’17. Red Hook, NY, USA: Curran Associates Inc., 2017, p. 4768–4777.
- [35] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine Learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [36] G. O. Campos, A. Zimek, J. Sander, R. J. Campello, B. Micenková, E. Schubert, I. Assent, and M. E. Houle, “On the evaluation of unsupervised outlier detection: measures, datasets, and an empirical study,” *Data Mining and Knowledge Discovery*, vol. 30, pp. 891–927, 2016.

Appendices

A Detailed analysis of the dataset

In this section, we present some additional insights on the LCPR dataset. Figure 1 presents the distribution count of some key features for each substation. We observe that meteorological features are identical for each substation as they are geographically adjacent and located in dense neighborhoods.

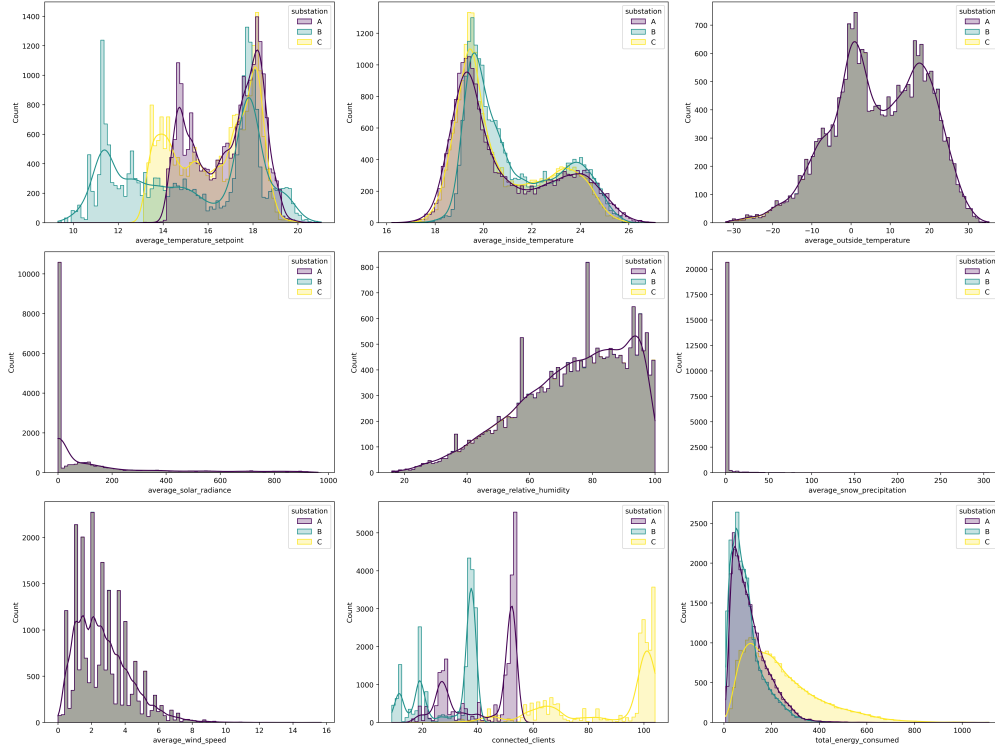


Figure 1: Distribution of key features for each substation

Figure 2 shows a correlation heatmap of important features and label for each substation. We observe that each substation follows the same general tendencies. We remark significant correlations between the energy consumed, the month of the year, the outside temperatures, and the temperature setpoints. This is also highlighted in Figure 3 which presents the Spearman coefficients ranking [32] between each feature and the label. A positive sign indicates that both the label and the feature grow or decrease in the same direction while a negative sign indicates an opposite direction. The coefficients are ranked in decreasing importance from left to right.

To have a more nuanced analysis of the contribution of each feature to the output, we also provide in Figure 4 an analysis of Shapley values of a trained extreme gradient boosting model (XGBoost) [33] for each substation. These analyses were realized with the Python package SHAP [34]. As we see, some lower-ranked features, viz., the challenge flags, sometimes have a strong impact on the model’s output even though their general impact is null.

B Numerical study of the Sliced-Wasserstein Filter

We first begin this section by showing the AD mechanism of the SW filter on a simple two-dimensional example. We generate three Gaussian distributions with different population sizes. As shown in Figure 5, the first distribution is the majority group, the second is the minority group, and the third represents clear statistical outliers. We now merge the three distributions into a single one to test our SW filter. We vary the radius of the Wasserstein ball to see how the filter behaves. The results are presented in Figure 6.

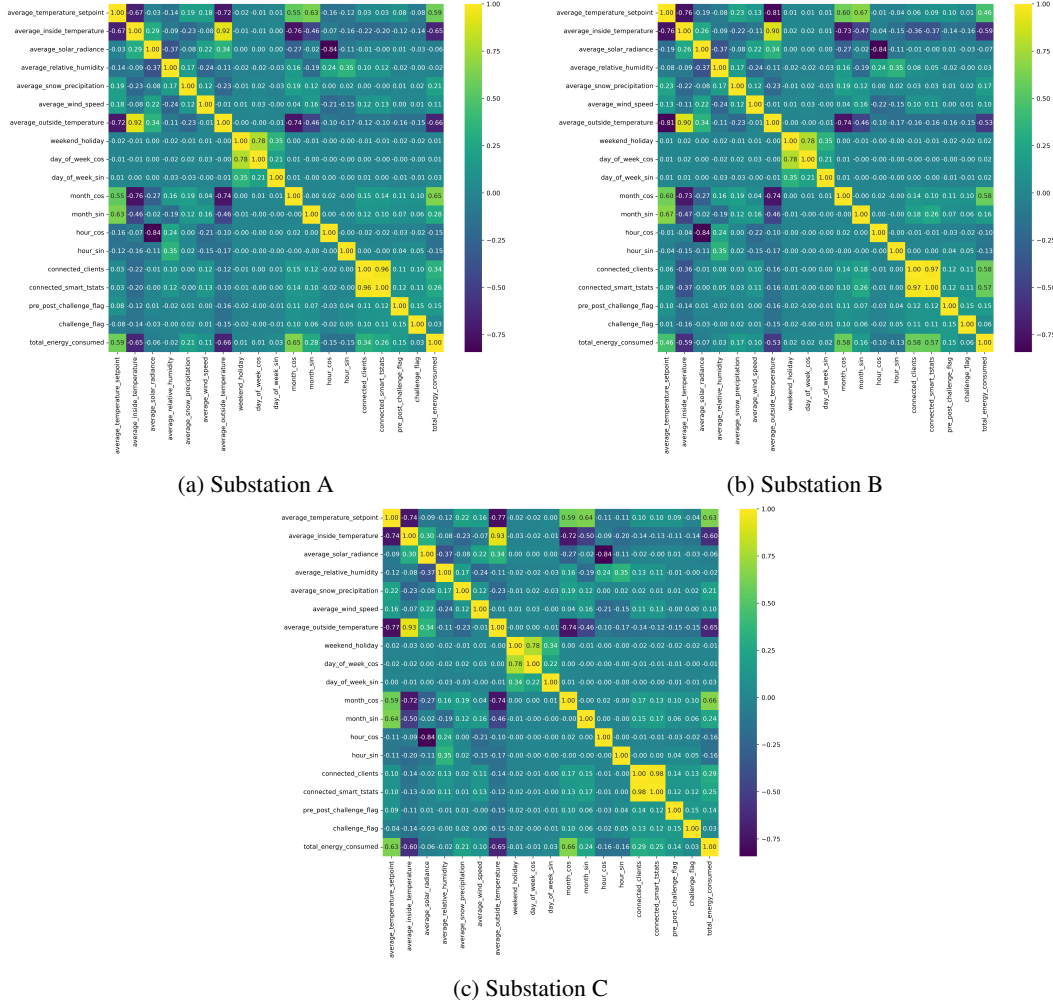


Figure 2: Correlation heatmap of key features and label for each substation

As we see, we can generate three filtering scenarios by modifying the value of ϵ . With $\epsilon = 0.1$, we filter only the statistical outliers. With $\epsilon = 0.05$, we only keep the majority group. And, with $\epsilon = 0.01$, we only keep the samples closest to the barycenter of the majority group. This is very interesting in a safe ML pipeline as we can tune the *conservatism* of the training dataset that is used at each run during hyperparameter optimization.

In Figure 7, we compare different AD algorithms on synthetic datasets provided in `scikit-learn`'s example collection [35]. Hyperparameters are fixed for each algorithm to see how a single hyperparameter choice influences the labelling on each dataset. As can be seen, the SW filter and its fast Euclidian approximation are better at isolating outliers when there is a clear majority group but are not as precise in identifying local outliers based on local density, e.g., one-class SVM.

Finally, we run a more thorough experiment with typically used real-world benchmark datasets. We run experiments on the *Lymphography*, *Ionosphere*, *Glass*, *Shuttle*, *WPBC*, *Arrhythmia*, and *Pima*, datasets presented in [36] for AD benchmarking. We compare isolation forest and LOF to our method. We implement a grid search and select each model's run with the best accuracy A . We then extract the precision score P from this run. The performance indicators are defined as follows:

$$P = \frac{T_p}{T_p + F_p}, \quad A = \frac{T_p + T_n}{T_p + F_p + T_n + F_n},$$

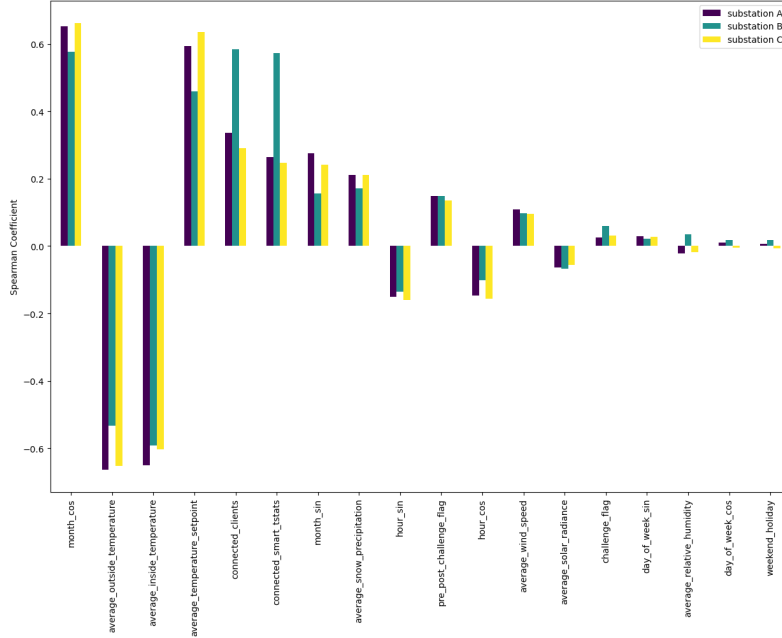


Figure 3: Spearman coefficients of key features for each substation

where T_p , F_p , T_n , and F_n stand for true positives, false positives, true negatives, and false negatives, respectively. Results are presented in Figure 8. We observe a similar performance between each model, except on the *Ionosphere* dataset where the SW filter lags behind the other models, and on *Arrhythmia* where the fast Euclidian approximation is underperforming. The SW filter’s strength is that it considers the global distributional properties of the population to guide its labelling. We remark that our method fails at detecting local outliers as it is purely designed to locate global outliers that are *risky* to keep in the training set. This is a reasonable choice to make when designing data preprocessing methods for safer ML training but not necessarily to filter local outliers.

Every experiment is available on our GitHub page.

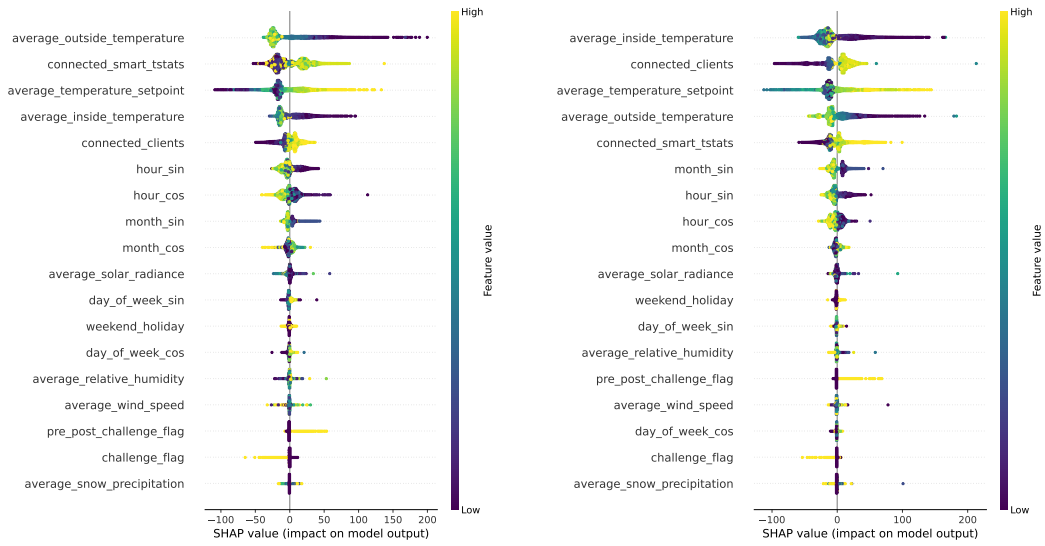
C Supplementary content to the benchmark

This section presents visual test predictions of the best validation runs for each substation as well as the test error of the benchmark.

As can be seen in both Figure 9 and Table 2, substation C is the hardest to predict for the Gaussian process. In general, the confidence of the model is also low (standard deviation is high). This hints at the complexity of the patterns.

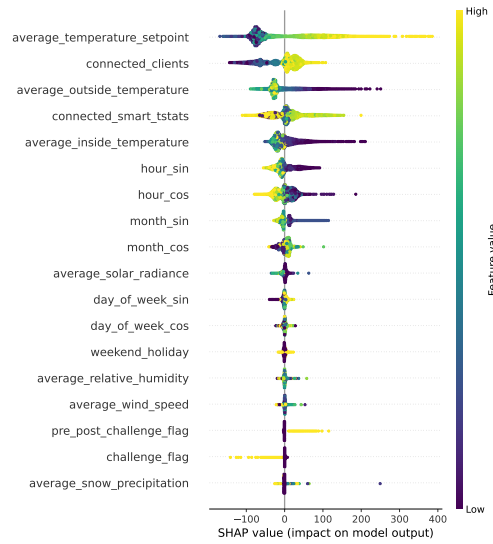
D Acknowledgements:

Special thanks to Odile Noël, Ahmed Abdellatif, Steve Boursiquot, and everyone from Hydro-Québec’s open data initiative, especially Andrée-Anne Gauthier, and Robert Row. This work was partially funded by NSERC, FRQNT, Mitacs, and Hydro-Québec.



(a) Substation A

(b) Substation B



(c) Substation C

Figure 4: Shapley analysis of key features for each substation on trained XGBoost

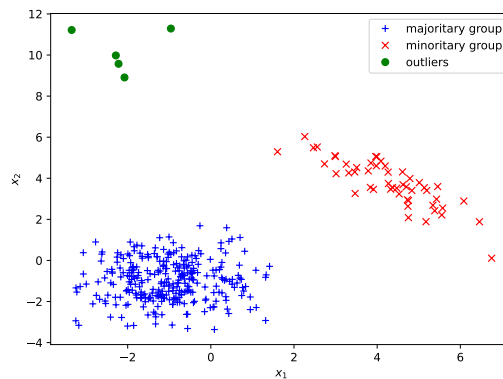


Figure 5: Illustration of different groups

Table 1: Description of features and label of the dataset

Name	Description	Possible values
substation	Substation identifier	{'A', 'B', 'C'}
timestamp_local	Timestamp in local time (UTC-5) and ISO 8601 format [AAAA-MM-DD hh:mm:ss]	–
connected_clients	Number of clients connected to the substation during the considered hour	{9, 10, ..., 104}
connected_smart_tstats	Number of smart thermostats connected to the substation during the considered hour	{59, 60, ..., 1278}
average_inside_temperature	Hourly average indoor temperature measured by smart thermostats in substation [°C]	[16.21, 27.08]
average_temperature_setpoint	Hourly average setpoint of smart thermostats in substation [°C]	[9.31, 21.03]
average_outside_temperature	Hourly average outside temperature at substation [°C]	[−32.0, 35.2]
average_solar_radiance	Hourly average solar radiance at substation [W/m ²]	[0, 961]
average_relative_humidity	Hourly average relative humidity at substation [%]	[0, 100]
average_snow_precipitation	Hourly average amount of snow precipitation at substation [mm]	[0.0, 306.0]
average_wind_speed	Hourly average wind speed at substation [m/s]	[0, 15.68]
date	Date [AAAA-MM-DD]	[2022-01-01, 2024-06-30]
month	Month	{1, 2, ..., 12}
day	Day of the month	{1, 2, ..., 31}
day_of_week	Day of the week with Sunday and Saturday being 1 and 7, respectively	{1, 2, ..., 7}
hour	Hour of the day	{0, 1, ..., 23}
challenge_type	Type of challenge during the given hour	{'None', 'CPR', 'LCPR'}
challenge_flag	Flag indicating hours in challenge	{0, 1}
pre_post_challenge_flag	Flag indicating hours in pre-challenge or post-challenge	{0, 1}
is_weekend	Flag indicating weekends	{0, 1}
is_holiday	Flag indicating Québec holidays	{0, 1}
weekend_holiday	Flag indicating whether a weekend or a holiday	{0, 1}
total_energy_consumed	Hourly energy consumption of the substation [kWh]	[7.45, 32240.17]

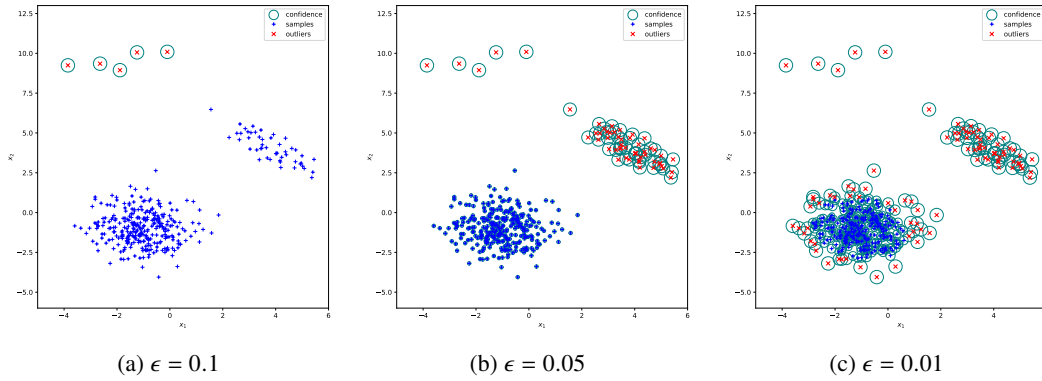


Figure 6: Labelling of the SW filter for different values of ϵ

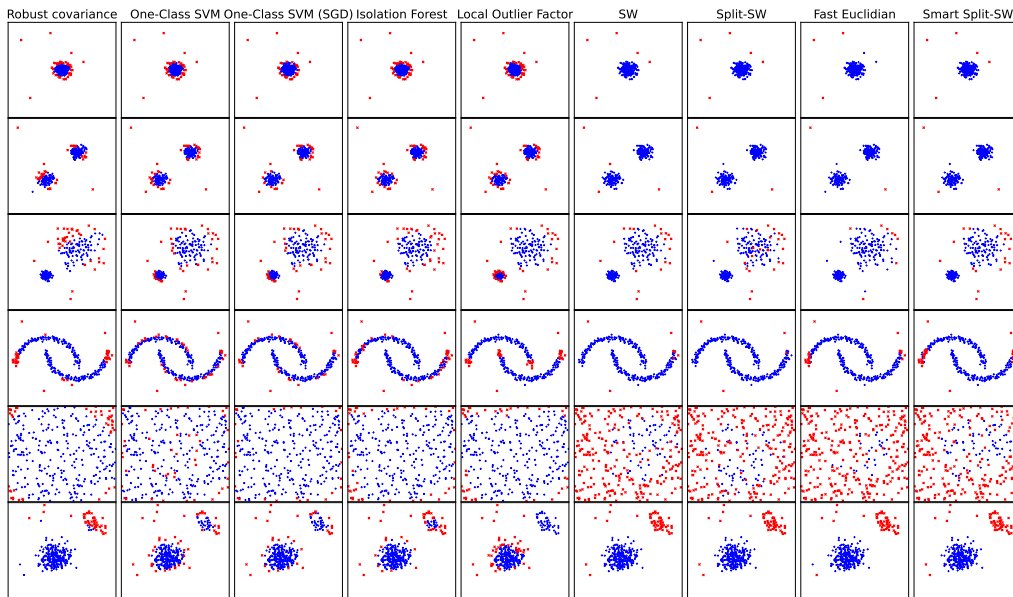
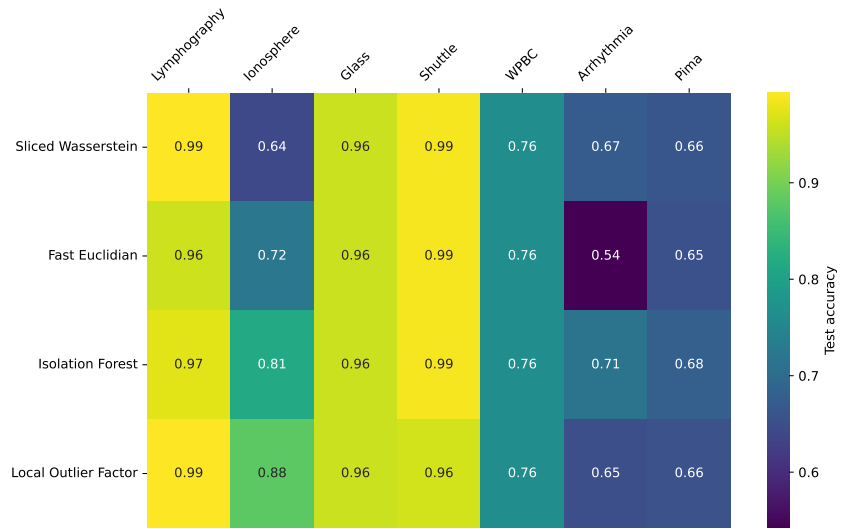


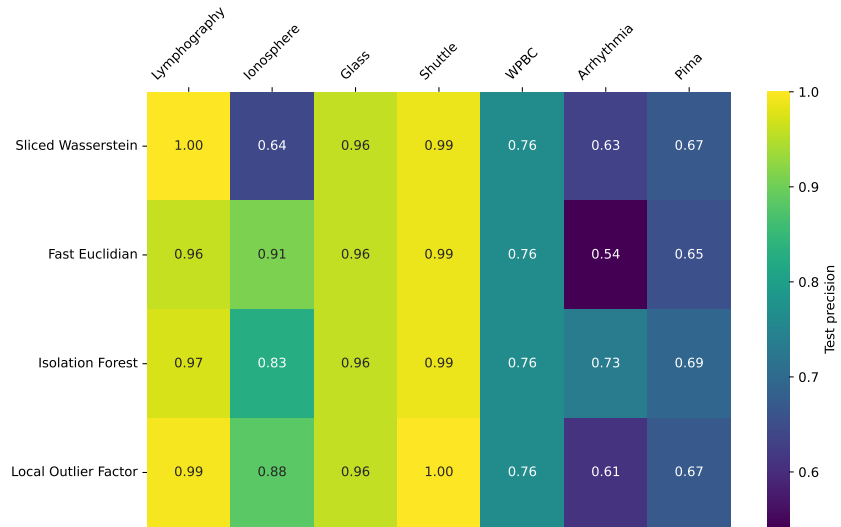
Figure 7: AD comparison on multiple synthetic datasets

Table 2: Absolute test errors of the best validation run for each substation

Substation	A	B	C
MAE [kWh]	20.49103	17.90663	41.21746
RMSE [kWh]	26.08270	22.39355	51.25044

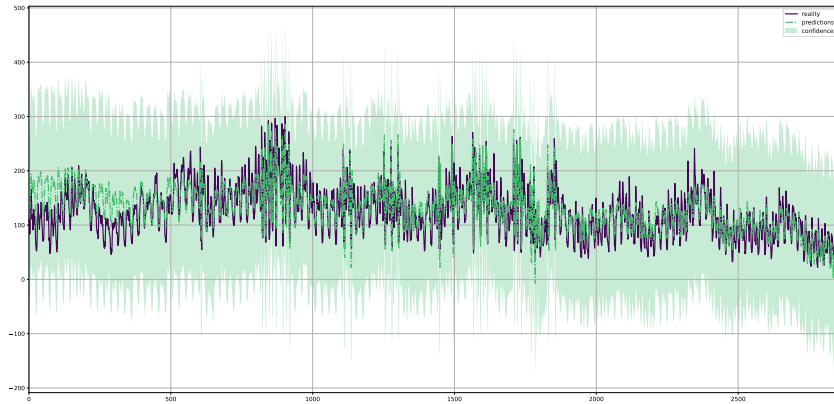


(a) Accuracy A

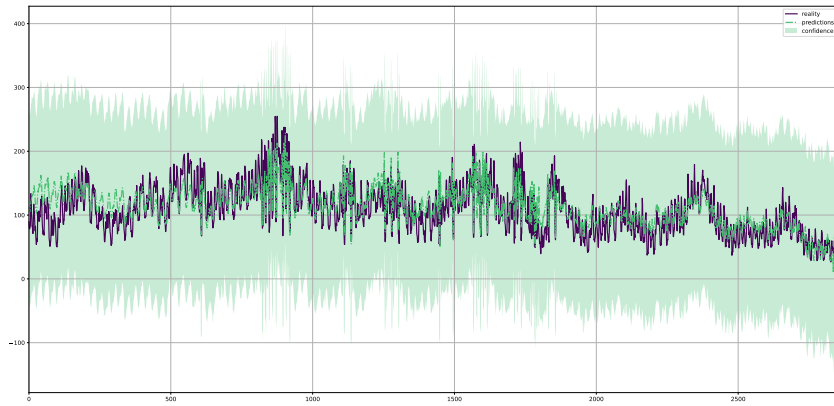


(b) Precision P

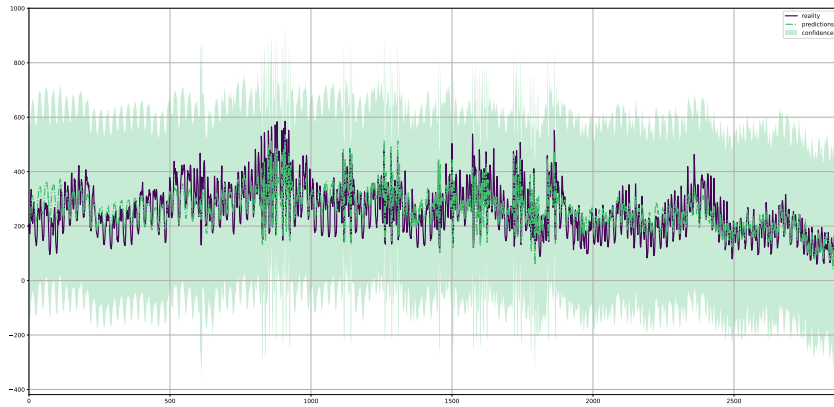
Figure 8: Results of the grid search for each AD model on each dataset



(a) Substation A



(b) Substation B



(c) Substation C

Figure 9: Test predictions of the benchmark at each substation