Overcoming Existing Limitations in Electricity-based Artificial Intelligence Applications

(Doctoral Consortium)

Mark Valovage
Computer Science and Engineering
University of Minnesota, Minneapolis, MN
valovage@cs.umn.edu

1. ELECTRICITY-BASED AGENTS

Machine-learning algorithms have recently been applied to electrical power problems due to their potential to reduce waste and improve electrical grid reliability, but deployment of existing research is hampered by unrealistic assumptions. My thesis focuses on overcoming these assumptions.

Autonomous agents have enormous potential to reduce electrical waste and improve reliability of the power grid. Imagine an 'energy feedback' agent at your home or business that identifies electrical waste in real-time, makes recommendations on when to operate different appliances, schedules your dishwasher to run during low-demand hours, and even recommends new appliances you can buy to replace your old, inefficient ones. Such an agent could eliminate billions of dollars of electrical waste every year. Similarly, picture an agent that monitors the power grid, identifies relay misoperations, and corrects them in a fraction of the time required by a human. This agent will have the ability to stop a localized misoperation from creating cascading events, preventing blackouts before they happen.

2. ELECTRICITY DISAGGREGATION

2.1 Label Correction

Electric companies bill consumers using aggregate power measurements. "You used X kilowatt-hours of power, pay us \$Y." In contrast, the goal of disaggregation is to separate individual appliances from the aggregate power signal and provide consumers with data for each appliance. This type of appliance-specific feedback has the potential to reduce power usage by an estimated 15% [3]. While this can be achieved through commercially-available smart meters, the expense of installing a smart meter for every appliance greatly exceeds the potential waste reduced. As such, both supervised and unsupervised learning methods have been explored to separate these appliance signals from building-level or circuit-level aggregate signals.

In supervised learning, the consumer has to capture isolated training samples for each appliance. Then, during normal usage, individual appliances have to be identified from the aggregate signal. Previous supervised learning research has assumed that samples provided by naive consumers are properly isolated and error-free, unrealistic assumptions that limit the application of existing classification methods.

We provided the first approach to automated label correction in electricity disaggregation in [5] using fixed increment adjustment. To correctly relabel appliance training samples, fixed increment adjustment modifies each ON and OFF label by a predetermined increment λ. It then generates each possible set of new labels and tests to ensure they meet two criteria. First, since the training sample must be captured in isolation, the power before and after the appliance operation must be the same (within a certain threshold ν). Second, the power increase during the sample must be positive and significant (higher than a predetermined threshold θ). Using these constraints enabled correction of 77% of training samples across four different houses and identified 2/3 of contaminated samples, but the method relied on the tuning of multiple parameters and required a pre-processing step of noise reduction as well.

To overcome these limitations, in [6] we introduced the application of Bayesian change detection (BCD) [1] to correct labels. BCD builds probabilities over all possible run sequences of observed data points. Because of the robustness of this probabilistic model, BCD showed similar performance to our previous label correction work, but did not require any pre-processing of the signal to reduce noise or tuning of specific parameters. We also showed this can be performed in real time with inexpensive hardware.

2.2 Event Detection

Unsupervised learning methods don't require training samples, eliminating hours of mundane setup time for the consumer. However, a lack of consumer input means unsupervised methods require at least two steps to discover what appliances exist in the monitored building. In the first step, referred to as event detection, the method must segment the raw, unlabeled power data stream by detecting significant electrical events. Existing event detection methods are heavily dependent on parameters pre-tuned for a specific dataset and sometimes even for specific appliances, techniques that do not scale to devices deployed to millions of buildings. We addressed this challenge in [6] by using the same BCD application described above.

The original BCD algorithm was designed to run online [1]. However, after an event occurs, online BCD requires a number of new observations before it can correctly identify the event. This can create a time delay between when

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

the event occurs and when online BCD detects it. Since BCD is probability-based, this time delay is even longer for lower power events. As such, we also introduced an offline BCD variant which computes all probabilities prior to event detection, eliminating the previously described time delays. Although this is done offline, near-online computation can be achieved by using a sliding window instead of processing the entire sequence of observed power values in batch [6].

2.3 Disaggregation Work in Progress

The second step in unsupervised electricity disaggregation is recombining detected events into their respective appliances. This is computationally challenging due to the sheer number of possible subsequence combinations. Even a very short sequence of 20 electrical events contains 2^{20} possible subsequences. Checking all of these individually is computationally intractable for any realistic sequence of events.

Previous work has addressed this obstacle in one of two ways. Brute force search can be used if the sequence length is capped to a very low number, such as 3. However, this only enables discovery of simple appliances. Alternatively, methods can start with generic models of appliances and attempt to match detected events to these. This can work on houses that have similar sets of appliances, but does not work for those that don't. In addition, appliances present in commercial buildings vary more widely than residential buildings, making generic appliance modeling infeasible.

Instead, I am currently working on a prioritized approach to appliance discovery that can intelligently explore the subsequence search space and does not require assumptions on the structure of appliances *a priori*. This appliance discovery only needs to be performed once (or possibly a few times, if new appliances are added to the building), so the goal of this research is to make appliance reconstruction computationally tractable using inexpensive hardware.

3. PMU EVENT CLASSIFICATION

Improving grid reliability is another promising application of machine-learning. Despite meticulous planning, relay misoperations (such as faults or loss of generation) occur throughout regional power grids. Misoperations can remain unnoticed due to the sheer size of transmission grids and a lack of global knowledge, and misoperations have historically caused or exacerbated large-scale blackouts [2]. A supervisory system could automatically correct for misoperations much faster than a human, but has been infeasible in the past due to limited data resolution. With the recent advent of phasor measurement units (PMUs), which can provide high-resolution, timestamped, and synchronized data about voltage magnitude and frequency data in the grid, such a system is realistic in the near future [2].

However, classifying misoperations needs to be done not just in real time, but with extremely high accuracy, as the entire purpose behind a supervisory system is to automatically correct such misoperations. If any event is misclassified by the supervisory system, it could initiate an incorrect action which could magnify the effects of a misoperation instead of compensating for it.

To meet the demands of rapid classification, time series shapelets have been explored as a possible classifier for PMU events. A shapelet [7] is a small portion of a time-series that can be used for classification. Shapelet discovery methods have provided high accuracy on dozens of other datasets [4, 7], but perform poorly when applied to PMU event data. Current state-of-the-art work relies instead on domain-level heuristic anchor points to extract shapelets capable of providing any level of accuracy [2]. I am currently exploring ways to modify shapelets beyond using raw distance classification metrics to make them robust enough to accurately classify PMU disturbance events. These improvements may improve shapelet accuracy in additional domains as well.

4. CONCLUSIONS

Although many supervised and unsupervised learning approaches have been explored for electricity disaggregation classification, their accuracy depends on unrealistic assumptions. A portion of samples provided by naive consumers will undoubtedly contain errors and some will be contaminated by other appliances [5, 6]. Widespread application of unsupervised learning is limited by individual parameters that must be tuned to a specific house or for specific appliances. Our modifications and application of BCD in [6] overcomes both of these limitations.

Similarly, shapelet discovery breaks down when attempting to identify key features in PMU event data. Applying a specialized agent that uses domain-specific constraints enables these methods to perform more realistically and effectively. While specific domain-level properties used may be unique to electrical data, other time series domains likely contain similar properties of their own that can be exploited to improve robustness and classification of applied methods.

REFERENCES

[1] R. P. Adams and D. J. MacKay. Bayesian online changepoint detection. *arXiv preprint arXiv:0710.3742*, 2007.

[2] M. Biswal, Y. Hao, P. Chen, S. Brahma, H. Cao, and P. De Leon. Signal features for classification of power system disturbances using PMU data. In *Power Systems Computation Conf.*, pages 1–7. IEEE, 2016.

[3] K. Carrie Armel, A. Gupta, G. Shrimali, and A. Albert. Is disaggregation the holy grail of energy efficiency? The case of electricity. *Energy Policy*, 2012.

[4] T. Rakthanmanon and E. Keogh. Fast shapelets: A scalable algorithm for discovering time series shapelets. In *Proc. of the 2013 SIAM International Conference on Data Mining*, pages 668–676. SIAM, 2013.

[5] M. Valovage and M. Gini. Automatic label correction and device prioritization in single household electricity disaggregation. *Workshop on Artificial Intelligence for Smart Grids and Smart Buildings, at AAAI*, 2016.

[6] M. Valovage and M. Gini. Label correction and event detection for electricity disaggregation. In *Proceedings of the 2017 International Conference on Autonomous Agents & Multiagent Systems*, 2017.

[7] L. Ye and E. Keogh. Time series shapelets: A new primitive for data mining. In *Proc. of the 15th ACM SIGKDD International Conf. on Knowledge Discovery and Data Mining*, pages 947–956. ACM, 2009.