21, rue d'Artois, F-75008 PARIS
http : //www.cigre.org

CIGRE US National Committee
2021 Grid of the Future Symposium

Machine Learning in Distribution Management Systems as Forecasting

S.P. MORASH[1], A.F. SNYDER[1], M. HUMAYUN[1], J. SCHOENE[1], Y. NAMAN[2]
EnerNex[1], Northeastern University[2]
USA

SUMMARY

This paper proposes a new machine learning framework applied to modern advanced distribution management system (ADMS) modules, including distribution system state estimation, distributed energy resource management, and fault location isolation and service restoration. The expectations of the ideal ADMS have seemingly continued to change with dual developments in machine learning techniques (including computational advances) and the transition to smart grids, which change the underlying data structure. Based on previous research, the framework assists in achieving the ambitious goals of modern distribution management systems by self-defining the level of granularity expected and refresh rate of the ML algorithm. One of the key findings of the research and the proposed approach is to maintain a forward-looking perspective in system operations rather than an enhanced accuracy of historical performance.

The ADMS is an operational tool that administers advanced applications through continual power system evaluation. The ADMS can also operate in "study" modes that allow users to evaluate the system under different hypothetical conditions and configurations. Many of these tools rely upon a distribution system state estimation (DSSE) solution to understand system conditions, enabling system operators to make decisions based on the status of the power system. Determining how often the DSSE should arrive at a new solution is usually bound by the computational capabilities available within a specified (and often incomplete) budget. For an industry charged with powering the modern economy, this type of optimization is unsatisfying.

An approach that could make more sense is to refresh the state estimation upon receipt of unexpected information. Rather than thinking of a DSSE as a look at the very recent historical performance of the distribution system, the DSSE should consider what the very near future of the distribution system will be. This mostly reflects the theory presented in [1], wherein Jeff Hawkins outlines the memory-prediction framework. The framework established the theory that the human brain is a feed-forward hierarchical state machine with special properties that enable it to learn. As a "feed-forward state machine", the machine (our brain) responds to future events predicted from past data. Hawkins makes the case convincingly by noting how startling it is for us, as humans, to trip and fall, or to reach for a door handle and miss. By constantly having expectations about our environment, we're able to quickly identify when one thing is out of line.

The work presented here identifies how Hawkins' research can be applied to advanced applications of an ADMS, helping to synthesize the growing set of available research in power system control with the goal to integrate more effective algorithms into utility practices. The feed-forward prediction paradigm is usually applied within the pursuit of strong artificial intelligence – machines capable of healing themselves and being self-aware. Certain advocates share similar ambitions for the electric

smorash@enernex.com

power system. By evaluating new research with a lens towards how it integrates with a feed-forward paradigm, practitioners will be able to both simplify the body of research and identify the clear goal of improved foresight, which underpins most utility decisions and is capable of making a difference in utility operational efficiency.

KEYWORDS

Machine learning, distribution management, advanced distribution management system, artificial intelligence, distribution state estimation, distributed energy resource management

INTRODUCTION

The modern distribution management system is tasked with controlling the complex medium and low voltage network of electric power delivery. This includes not only the algorithms required to understand and manage the power flow under normal conditions, but to make sense of abnormal conditions, identify causes of abnormalities, and coordinate the repair, whether automated, remote, or requiring dispatch of a service crew to fix/replace broken equipment. The tasks of these modern advanced distribution management systems (ADMS) are critical to serve safe, reliable power to homes and businesses, the very mission of electric utilities. Because the ADMS combines analysis of complex networks of information with applications that must make that information useful and actionable for grid operators and service workers in real time, machine learning (ML) techniques can be useful.

Though the computer science is still maturing, machine learning has long been the topic of research for many individual components (modules) of modern ADMS solutions. Published in 1991, [2] builds from previous work to develop an artificial neural network to identify faults in advanced manufacturing facilities. Since then, neural networks have advanced from merely identifying fault conditions to predicting fault location in other sectors, including the work done in [3] to study the low voltage electric distribution network. Reference [4] developed a modified machine learning algorithm to solve a slightly different problem: restoring power after a fault.

To effectively understand and predict outage causes and locations, a connectivity model is required. This connectivity model effectively traces what equipment is located where and how it is all joined to deliver power. Utilities often struggle to assemble the connectivity model, which outlines electronically where and how different points on their system are connected. Reference [5] outlines a few approaches to how machine learning approaches have been applied to building the network topology configuration. For instance, [6] combines many statistical strategies and relies on smart meter data to reconstruct the partially or fully unknown lightly meshed topology. Accuracy is shown to be dependent on the number of measurement samples and complexity of system, but provides acceptable performance under a variety of conditions, including with missing nodal data. As outlined in [7], maintaining a highly accurate connectivity model is difficult in practice due to the dynamic nature of distribution systems, and insufficient circuit-level data to monitor topology changes.

Recognizing that more nuance is required to fully understand the power system than connectivity, [8], [9] and [10] have focused on developing machine learning techniques for distribution power flow and distribution state estimation. The use cases for these applications are numerous and cover a spectrum of distribution operation, including real-time situational awareness and identifying topological errors in the underlying connectivity model.

There are many more journal articles outlining innovative machine learning techniques applied to the complex distribution system challenges facing the utility of the future. However, [11] outlines many of the factors inhibiting machine learning adoption in the utility sector, including tactical challenges like how underlying data structures can be noisy and inconsistent, or that data security concerns can be misaligned with the massive computational burden required by some solutions. The article goes on to identify that many decision-makers do not understand the underlying mathematics, nor are they willing to entrust key decisions to an opaque algorithm.

Even more fundamentally, this paper posits that machine learning applications in the utility space have failed to deploy rapidly because they lack an overarching framework within which to align. Vendors and researchers deploy machine learning techniques to tackle specific problems, aiming today's robust computational capabilities at subsets of the challenges, seemingly embracing a weak AI approach. To combat this and to make better use of the capabilities available thanks to decades of research in both machine learning and distribution system management, the concept of a feed-forward state estimator is introduced in the next section before exploring how that feed-forward engine would impact related tasks facing modern electric power management. Finally, next steps will be discussed.

A MACHINE LEARNING PREDICTION ENGINE FOR ADMS

The ADMS is an operational tool that administers advanced applications through continual power system evaluation. The ADMS can also operate in "study" modes that allow users to evaluate the system under different hypothetical conditions and configurations. Many of these advanced applications rely upon a distribution system state estimation (DSSE) solution to understand system conditions, enabling system operators to make decisions based on the status of the power system. The authors of [12] outline the imperative nature of this state estimation before introducing an algorithm to solve the task. However, they do not cover how often the algorithm should be run nor how often circuit-level telemetry should be scanned (polling frequency). Additional research found minimal published industry input on the optimal DSSE execution interval and polling frequency. Determining how often the DSSE should arrive at a new solution is partly bounded by the computational capabilities available within a specified (and often incomplete) budget and also by the polling frequency, which is constrained by sensor capabilities and the communication infrastructure. For an industry charged with powering the modern economy, this type of optimization is unsatisfying.

An approach that addresses the limits due to computational capabilities is to refresh the state estimation by exception, i.e., upon receipt of unexpected information. Rather than thinking of a DSSE as a look at the very recent historical performance of the distribution system, the DSSE should consider what the very near future of the distribution system will be. This mostly reflects the theory presented in [1], wherein Jeff Hawkins outlines the memory-prediction framework. The framework established the theory that the human brain is a feed-forward hierarchical state machine with special properties that enable it to learn. As a "feed-forward state machine", the machine (our brain) responds to future events predicted from past data. Hawkins makes the case convincingly by noting how startling it is for us, as humans, to trip and fall, or to reach for a door handle and miss. By constantly having expectations about our environment, we're able to quickly identify when expectations are not met.

Bringing this idea back to DSSE and the ADMS, most system operations assume that the near future will look very similar to the recent past. Much literature is focused on more accurately understanding what just happened, including [13] and [14]. However, by looking forward, the system can more readily identify potentially troublesome deviations in performance. A near-term prediction would help the system understand that new observations are within the realm of the expected; it will more readily identify when things aren't "normal."

Mostly, this prediction engine focuses on a DSSE tolerance for error as a refresh function rather than time. This is key to the entire framework. The feed-forward perspective feeds from historical data, aiming to predict the future, but also correcting the predictions based on new information as necessary. By staying consistently ahead of the current moment, the DSSE is afforded more time to generate predictions before they are absolutely required. The framework also allows for evaluation of individual sensor measurements relative to expectations as they arrive without running a full DSSE. Real-time measurements affect the future estimation at different scales and could be incorporated as the error tolerance dictates.

The feed-forward framework is not a substitute for machine learning. Machine learning can be applied in post-event processing. Rather, the framework is a lens through which to evaluate operational machine learning. Quite simply, the framework starts with the following fundamental questions:

- Does the machine learning application help the electric system respond to future events?
- If so, does it do it in a useful fashion[1]?

This feed-forward DSSE framework is preferable because of real world constraints. Of course, equipping systems with the latest information will always make a forecast better. Of course, building a new estimation based on every new piece of information is ideal. However, the pursuit of fast, accurate,

[1] Examples of applications that are not useful include those that are: too expensive, too slow, or too sensitive to bad data.

and resilient DSSE has been the topic of much research, but with few of these solutions migrating from research to commercially available ADMSs. Reference [7] posits that "data scarcity inhibits the adoption of advanced DSSE techniques into today's ADMSs and today's ADMSs are not capable of fully leveraging circuit-level data…" However, this paper builds upon that underlying, data scarcity problem with an additional, architectural one: industry must know when to stop; when perfect forecasting, state estimation, or topological error detection should be secondary to enabling advanced applications. Effectively, research and vendor development must prioritize approaches that are feasible, scalable, affordable, and capable of unlocking value on the short term while the underlying data problems persist.

By focusing on the near-future only, deviances from perfect accuracy (from a perfect electrical model) are expected. Forecasts are usually wrong, but are helpful, nonetheless. This approach is also more to the core of how the distribution system is operated. System operators do not need a DSSE with 0.2% accuracy at all nodes: meters can have errors larger than that, voltage standards allow for 5% shifts in either direction from nominal [15]. The work here reflects the reality of many distribution system operators: they just need to know when the system is not performing as expected. This is not to say that a state estimation should only be refreshed under some alarm condition, but rather that it should mirror the dynamic nature of the physical system that it models.

The feed-forward framework – effectively an emphasis on proactive forecasting for distribution system management- should lower the computational burden placed on operational systems. The system does not need to constantly fit each new piece of information into a new state estimation. Rather, it assesses the information, determines its validity, aggregates that information, and uses it only when it becomes necessary (when the prediction is too wrong, or enough data is out of line with the forecast). Plenty of work has been done in short term forecasting that can be built upon and adopted. Reference [16] identifies the efficacy of neural networks and neuro-fuzzy algorithms to predict near-term voltages, ultimately determining that satisfactory results to voltage forecast can be made from 1-minute interval data looking 10 minutes ahead. Reference [17] outlines a neural network ensemble that predicts load at both a customer and substation level 24 hours ahead in a real-world setting. Countless other journal articles focus on predicting load, particularly in environments with behind-the-meter solar generation. All of this will be needed to adopt a feed-forward paradigm.

The work presented here differs from other approaches in that the emphasis is not accurate, *real-time* status of the distribution system, but orients both researchers and real-world distribution operations around *near-future* status. In some places, the system operator is expected to carry the burden of looking into the future. In other instances, ADMS's utilize short-term forecasts to recommend system adjustments. These architectures place a burden on other systems to interpret the current status of the system and make informed projections from it. The feed-forward architecture orients every module around a common operational time domain with each module containing running their requisite granular algorithms to make informed decisions. Effectively, real-time information and status are too slow to take appropriate proactive/preventative measures. System operators really need a look into the future to head off potential problems. It makes little sense to be concerned with the accuracy of the current state, other than to contextualize our projections.

FEED-FORWARD ARCHITECTURE FOR DER MANAGEMENT

As covered earlier, our electricity system is shifting to enabling active customer participation in the grid with growing importance placed on accurately predicting end-user electricity profiles, particularly with distributed generation. Affecting change within this segment requires advance notice: these resources are not typically integrated within traditional, automated generation control architectures and do not figure to be so soon. Managing DERs requires a prediction engine today and figures to be more reliant upon that feed-forward framework in the near future.

Dispatching DER

The time domain for distributed resource dispatch, affecting the status of the grid, must be forward looking. Resources need to know what is expected of them and given time to perform. DER programs, such as those being implemented by Hawaiian Electric Company [18] and elsewhere, outline the market procedures and performance expectations of DER when given varying forewarning. At their core, DER programs require foresight into grid conditions, communication to the DER of the expected remedial action, and performance of the DER. Depending on the program, the prediction engine and foresight required for active DER programs may vary from minutes to hours/days into the future. Enabling such programs, which figure to be more prevalent soon, requires an architecture oriented around a prediction engine to afford adequate time to influence customer behavior with the appropriate temporal and geospatial fidelity. Reference [4] identified this as an area in need of additional research. More specifically, the authors were interested in how price-sensitivity and/or prosumers impact DSSE.

Overcoming Dispatch Errors

Relying on DER for control of the grid comes with unique challenges, partially outlined within [19], that include the reliability of resources to precisely and accurately perform as dispatched. On top of that, there are challenges with evaluating the performance of DER, particularly when dedicated metering is unavailable. These operational challenges could be partially remedied with a feed-forward architecture that expects errors and remains dynamic in its dispatch by continuing to be forward looking. Dispatching with enough advance notice allows for additional dispatch as performance deviates from dispatch instructions. More directly, the real-time architecture does not care why or even that there are errors, but merely requires a far enough lens into the future in order to do something about those errors before they affect real-time operation. By establishing state estimation refresh and now the dispatch of DER as error functions, rather than as time functions, one can more accurately maintain the system unbeholden to arbitrary time increments.

Put into a graphical format, the forward-looking dispatch of DER might look like the Figure below. Note that the time domains used here are for discussion purposes only and may vary significantly depending on the feed-forward implementation.

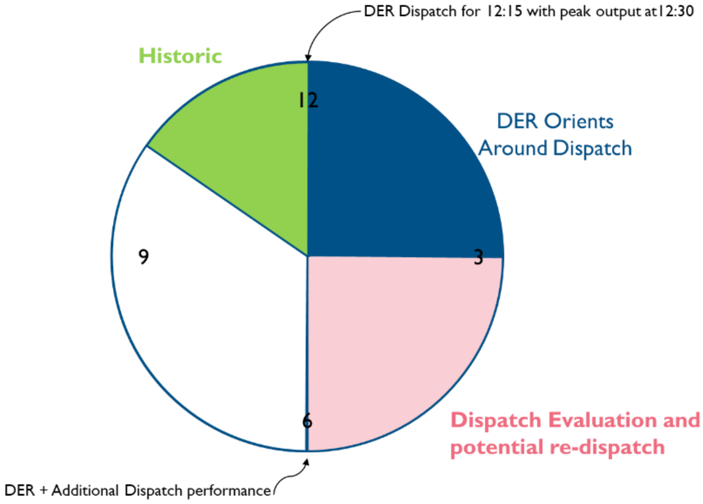

Figure 1 Multiple DER Dispatch Paradigm

This paradigm is not particularly novel; electricity markets perform these functions today. However, this discussion is included in this paper to reinforce the notion that *if* the future of our electrical system requires DER, then it requires a *forward-looking* perspective on the status of the grid at multiple temporal and geospatial resolutions. Effective DER dispatch cannot be done once, nor in real-time. It requires advanced planning with foresight ranging from years to minutes, in a concerted effort. Certain DER customers will opt-out, regardless of the financial incentive(s) provided. Just as with today, generators will be needed to balance any deficiencies caused by their peers. Certain DER customers will opt-out, regardless of the financial incentive(s) provided, while others will perform at a level greater than dispatched. Additional resources will need to balance the deficiencies in other, previously dispatched resources in order to maintain balance.

FEED-FORWARD ARCHITECTURE FOR UTILITY CONTROLLED EQUIPMENT

Whereas Section C focuses on market mechanisms, this section will target utility-controlled equipment, whether in fault location, isolation, and service restoration (FLISR), or in longer-term asset management.

Efficacy of Feed-Forward Paradigm in Outage Conditions

There are a variety of outage conditions and each has volumes of research dedicated to a more effective response. The purpose of this section is to simply discuss outage response as it relates to a feed-forward, prediction paradigm. At least one commercially available outage management system considers itself a “prediction engine” meaning that it analyzes available information and provides operators with a likely fault location and failed equipment. A strong machine learning approach combined with a perfect prediction engine would limit outage conditions by avoiding abnormal conditions. However, even a strong model of the physical world figures to be unable to predict squirrel behaviour in the foreseeable future, which unfortunately will result in outages.

Within [3], the authors outline a procedure that reduces the computational burden in identifying outage location by reducing the search space according to time series metering information. They then perform an automated voltage profile analysis using a Monte-Carlo algorithm. This method of predicting the outage location and then refining it with further analysis is well aligned with both today’s utility practices and the feed-forward architecture.

In fact, it seems that some of the more advanced machine learning research could benefit from similarly narrowing their algorithmic solution domain. The authors of [4] present a modified Q-learning method that simultaneously restores service and manages load under stressed grid conditions. The authors point out that one disadvantage of the solution is that “the utility should have complete information for the grid”. That documented disadvantage, the computational burden, and ultimately the time to run the solution could benefit from an architecture more reliant upon a prediction engine to narrow the scope of analysis.

Asset Management

Reference [20] focuses on a long-term asset health framework, building a machine learning algorithm to “construct a predictive model to determine the effect of weather events on network transformer longevity.” A feed-forward approach to this capability would transition the model from evaluating transformer likelihood to survive a given time frame to a prediction engine evaluating asset failure likelihoods given a particular set of near-future weather conditions. Such conditions may include an approaching hurricane, or extreme cold/hot. A dynamic understanding of asset performance under extreme weather conditions could help operators to more quickly locate and isolate failed equipment by narrowing the list of likely problem areas.

The feed-forward architecture for asset management does not just build upon weather events, it would also include asset management approaches when identifying and recommending switching procedures. Just as with today, when ADMSs evaluate thermal and other limits for distribution switching configurations, the feed-forward architecture would consider the lifespan of the assets prior to executing a new topology.

CONCLUSION

The work presented here identified machine learning research applied to advanced applications of an ADMS. Each research project contributed to a growing set of available, improved methods for integration within utility practices. Among other reasons, the overwhelming nature of the continuously growing body of machine learning options available has led to the slow adoption of these solutions in the marketplace. Within this paper, a framework through which to evaluate new machine learning solutions has been presented. The feed-forward prediction paradigm is usually applied within the pursuit of strong artificial intelligence – machines capable of healing themselves and being self-aware. Certain advocates share similar ambitions for the electric power system. By evaluating new research with a lens towards how it integrates with a feed-forward paradigm, practitioners will be able to both simplify the body of research and identify the clear goal of improved foresight, which underpins most decisions and is capable of making a difference in utility operational efficiency.

This framework borrows from existing research and certain portions have already been applied successfully to some use cases. Just as with other frameworks, its usefulness will be apparent when applied to a replicable demonstration project, one that perhaps tests the edges of what is possible.

One of the goals of the framework is to elevate the machine learning literacy and vocabulary of the utility workforce. The research and models discussed herein contain complex mathematical models that attempt to more effectively do relatively simple jobs, such as identifying the failed equipment causing an outage. Improving the collective understanding of machine learning in an industry starts with effective communication of the solution but is assisted by the framework presented herein. In order to understand something complex, it’s helpful to break it into smaller bites.

BIBLIOGRAPHY

[1] T. Sorsa, "Neural Networks in Process Fault Diagnosis," *IEEE Transactions on Systems, Man, and Cybernetics,* vol. 21, no. 4, pp. 815-825, 1991.

[2] "Data-Driven Fault Location Scheme for Advanced Distribution Management Systems," *IEEE TRANSACTIONS ON SMART GRID,* vol. 10, no. 5, pp. 5386-5396, 2019.

[3] L. R. Lambert-Torred, A. R. Ferreira and G. Aoki, "A Reinforcement Learning Approach to Solve Service Restoration and Load Management Simultaneously for Distribution Networks," *IEEE Access,* 2019.

[4] K. Dehghanpour, Z. Wang, J. Wang, Y. Yuan and F. Bu, "A Survey on State Estimation Techniques and Challenges in Smart Distribution Systems," *IEEE Transactions on Smart Grid ,* vol. 10, no. 2, pp. 2312-2322, 2019.

[5] D. Deka and e. all., "Structure Learning in Power Distribution Networks," *IEEE Transactions on Control of Network Systems* , vol. 5, no. 3, pp. 1061-1074, 2018.

[6] J. Schoene, M. Humayun and a. et, "Quantifying Performance of Distribution System State Estimators in Supporting Advanced Applications," *IEEE Open Access Journal of Power and Energy,* vol. 7, pp. 151-162, 2020.

[7] L. Yuxiao, "Data-Driven Power Flow Linearization: A Regression Approach," *IEEE TRANSACTIONS ON SMART GRID,* vol. 10, no. 3, pp. 2569-2580, May 2019.

[8] A. &. S. N. Zamzam, "Physics-Aware Neural Networks for Distribution System State Estimation," 2019.

[9] L. Wang, *Physics-Guided Deep Learning for Power System Sate Estimation,* Tampa: University of Central Florida, 2019.

[10] H. K. Trabish, "How does AI improve grid performance? No one fully understands and that's limiting its use," Industry Dive | Utility Dive, 14 November 2019. [Online]. Available: https://www.utilitydive.com/news/how-does-ai-improve-grid-performance-no-one-fully-understands-and-thats-l/566997/. [Accessed 22 March 2020].

[11] A. Arefi and e. all, "An Efficient DSE Using Conditional Multivariate Complex Gaussian Distribution," *IEEE Transactions on Smart Grid,* vol. 6, no. 4, pp. 2147-2156, 2015.

[12] J. Hawkins, On Intelligence, United States: Times Books, 2004.

[13] X. Hu, H. Hu, S. Verma and Z.-L. Zhang, *Physics-Guided Deep Neural Networks for PowerFlow Analysis,* 2020.

[14] A. S. Zamam, X. Fu and N. D. Sidiropoulos, ""Data-Driven Learning-Based Optimization for Distribution System State Estimation,"," *IEEE Transactions on Power Systems,* vol. 34, no. 6, pp. 4796-4805, 2019.

[15] ANSI C84.1-2016.

[16] E. D. Garcia and e. all, "New Alternatives to Improve Advanced Distribution Management Systems Using Very Short-term Voltage Prediction," in *Proceedings of the 2014 49th International Universities Power Engineering Conference (UPEC)*, 2014.

[17] M.Saviozzi, S.Massucco and F.Silvestro, "Implementation of Advanced Functionalities for Distribution Management Systems: Load Forecasting and Modeling through Artificial Neural Networks Ensembles," *Electric Power Systems Research*, vol. 167, no. February, pp. 230-239, 2019.

[18] Hawaiian Electric Company, "Demand Response," July 2020. [Online]. Available: https://www.hawaiianelectric.com/products-and-services/demand-response. [Accessed July 2020].

[19] S. Morash, "Distributed Battery Energy Storage: Intro to Battery DR and How Baselining Techniques Can Fail," EE Online, 21 November 2017. [Online]. Available: https://electricenergyonline.com/energy/magazine/1065/article/Distributed-Battery-Energy-

Storage-Intro-to-Battery-DR-and-How-Baselining-Techniques-Can-Fail.htm. [Accessed July 2020].

[20] S. McCormick, "Using Machine Learning to Quantify the Impact of Weather on Transformer Failure Risk," in *CIGRE Grid of the Future*, Atlanta, 2019.