# Intelligent Outside Plant Power Operations with Machine Learning

A Technical Paper prepared for SCTE by

**Matthew Stehman**
Principal Machine Learning Engineer
Comcast
1800 Arch St, Philadelphia, PA
Matthew_Stehman@comcast.com

**Chris D'Andrea**
Data Scientist
Comcast
1800 Arch St, Philadelphia, PA
Christopher_D'andrea@comcast.com

**Ilana Weinstein**
Data Scientist
Comcast
1800 Arch St, Philadelphia, PA
ilana_weinstein@comcast.com

 

# Table of Contents



## List of Figures



## List of Tables

# 1. Introduction

As cable providers continue to develop and deploy 10G technologies, customer experience expectations are also growing in stride. These more reliable networks and faster speeds mean that even the smallest outages have a high impact on daily life. In the realm of outside plant power, the power supplies are responsible for powering the plant as well as providing back up battery power and keeping the cable plant online during commercial power outages (CPOs). Artificial intelligence (AI) can be integrated with real-time telemetry to plan resource deployments based on the expected performance of the batteries combined with predictions on customer impacts during the outage. While CPOs are likely unavoidable for multiple systems operators (MSOs), advanced techniques can be employed to mitigate potential connectivity outages during such events.

The authors will discuss Comcast's approach to minimizing customer impact during CPOs using Artificial Intelligence and Machine Learning (ML). Big Data tools and pipelines are employed to gather power supply telemetry and customer connectivity data for historical CPOs. Predictive models are built to learn from these historical events and can provide insights during active CPOs. These predictive models can then be integrated into field operations platforms to help plan work and even proactively message customers. This paper will illustrate the power of using historical power outage data with AI/ML frameworks to support the connectivity expectations of a 10G network.

# 2. The Life Cycle of a Commercial Power Outage

The feeling during the first few minutes of a power outage is all too familiar to most people; lights go out, a scramble to find a light source begins (candle, flashlight or more likely a smart phone) and then the discussion of "how long will this one last." During commercial power outages, impacted customers are focused on what the power company is doing to bring power back up, but few know how stressful these outages can be for MSOs as well. MSOs generally would like to avoid the scenario where customers find their cable services unavailable after commercial power is restored.

While MSOs are typically in a reactive state of operations during CPOs, smart data collection and predictive models can significantly help with more targeted operations and clearer communications to subscribers. The following sections will discuss the different stages of a CPO and how Comcast is using data-driven approaches to build AI/ML-driven systems to learn from real outage data.

## 2.1. Outage Detection

The first step in resolving a CPO is knowing there is one. The detection of a CPO starts with data over cable service interface specifications (DOCSIS) customer premise equipment (CPE) registration data. Typically, during a CPO, a cluster of modems will become unreachable around the same moment in time. The CPE polling systems will label these devices with offline statuses, and they typically remain offline together until services are restored. Once a group of devices is determined to be offline, a clustering algorithm is invoked to determine if they are part of the same outage. The next step is to rule out network plant damage, such as a fiber cut, or something related to a hybrid fiber coax (HFC) outage. Comcast's HFC network is built upon a vast Outside Plant (OSP) power network with real-time telemetry. If MSOs have visibility into CPE connectivity and OSP power supply statuses, the two can be correlated together to see if the group of offline modems have any corresponding power supply (PS) that also lost input power around the same time.

It is important to note that the OSP power supplies typically have battery backup systems to keep the HFC plant operational for a prescribed amount of time during typical CPOs. More information on the types of HFC-powering options during CPOs can be found in Peck and Frankhouser (2021). In these cases, there are two types of scenarios for CPE status: customers that do not have an external power source will register as offline, and customers that do have an external power source (home generator or local backup batteries) will remain online as long as the HFC plant is online or until their power source depletes. In these cases, the combination of scenarios needs to be considered as not all customers experiencing a power outage will appear offline immediately.

All previous steps can be automated using internal data sources. The final verification of a CPO can be made by cross-checking with the local power provider in the area of the outage. This can be done with automation and manual engagements depending on the circumstances. The final step can take some time as there may be different lag times between the various systems that need to communicate. It is thus the case that the automation should first create the outage as soon as the confidence level is high enough, and then when the CPO classification is made, the outage is updated, and a potentially different workflow is engaged.

A schematic representation of this process is described in Figure 1. The next section will discuss what actions are typically taken once the CPO is declared.

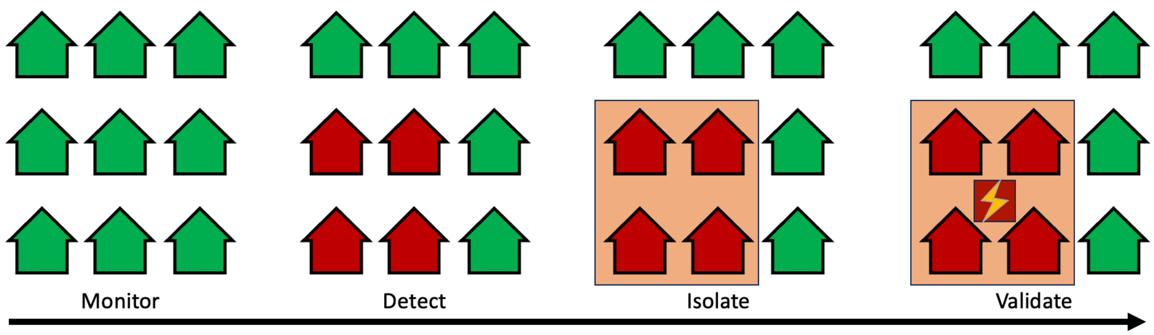


**Figure 1 - Detecting a CPO**

### 2.2. Workforce Planning and Maintenance Activity

Once the outage is declared the field teams take over and determine the necessary courses of action. An evaluation of the outage is usually the first task. This typically involves answering questions around: how many customers are offline, how many nodes are impacted, what's the status of our power supplies, is there an Estimated Time to Repair (ETR) from the power provider, is it weather related or something else. In this phase, teams are triaging the situation and getting as much information gathered as possible to determine what operations will be needed to resolve any issues. This data is extremely useful in cases where the demands of the outages exceed the available resources. In such cases, workforce planning efforts may be needed to determine the best sequence of operations based on a variety of factors.

Customer impact is at the forefront of the evaluation. Who is impacted and who may become impacted, may drive the maintenance activities in compliance with the Telecommunications Service Priority (TSP) program. To answer the first part of the question, who is impacted, the status of the CPE and power supplies give an initial estimate of the blast radius of the outage. To answer the second question, who may become impacted, there are two factors. The first is the progression of the outage, which is likely due to factors outside of the MSO's control. The second aspect is related to the power supply batteries mentioned in Section 2.1. During the loss of commercial input power, the OSP power supplies can run on battery to keep the HFC plant operational. However, the batteries have a finite run time. If the duration of the CPO exceeds the run time of the backup batteries the HFC plant will go down, resulting in customers being disconnected.

The potential impact to currently online customers' needs to be considered when planning resources and operational activities. Comcast employs a machine learning battery run time model to predict the run time of the power supply batteries at the beginning of every CPO. This run time prediction is then used to determine the need to roll out backup generators to support the HFC plant during outages that exceed the run time capacity of the power supply batteries. Since the run time prediction is made at the beginning of the outage, these generator decisions can be discussed very early in the planning process. Section 3.2 will discuss the run time prediction model in more detail. Figure 2 shows a schematic of the general information that should be considered when planning resources and potential actions to address a CPO.

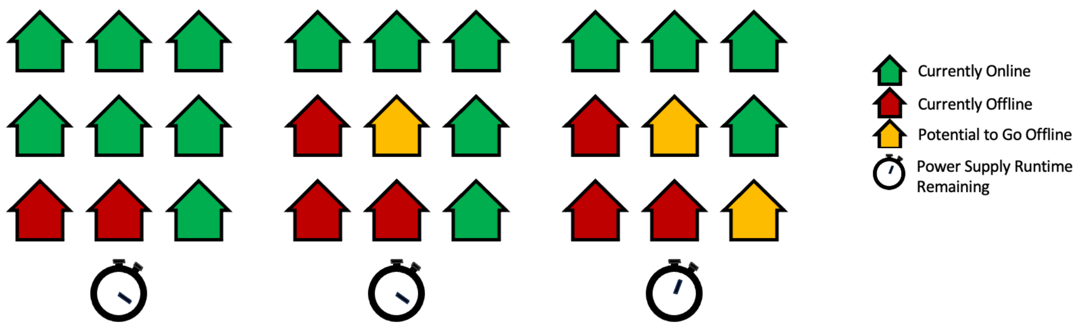


**Figure 2 - CPO impact assesments**

Another key aspect of the CPO action plan is customer communications. As shown in Figure 2, customers can be in a variety of states during CPOs and a message from their cable provider acknowledging the situation and explaining the next steps can go a long way. Pertinent information for these communications includes current status, possible future statuses, and timelines where available. These items can be derived by combining real-time telemetry with AI/ML systems that learn from historical outage data.

### 2.3. Outage Resolution and Closure

Once the outage begins to resolve and customers begin to come back online, procedures can be taken to start clearing the outages. Typically, this begins to show via telemetry in the form of the outage footprint shrinking as more customers begin connecting to the HFC plant and their DOCSIS registration statuses are online. Depending on the nature of the outage and the amount of time it takes all customers to come back online, a physical plant inspection may be needed. When customers are back online, the outage will be closed, and this lifecycle will begin again in the next outage.

When the outage is closed, some amount of data will need to be saved for documentation and outage history tracking (depending on the details of the outage, reports may be filed with the Federal Communications Commission and some state regulatory agencies). In addition, if machine learning techniques are to be employed to learn from historical data, as is done at Comcast, then some thought should be put into what data should be stored to compile a training data set for building predictive ML/AI models to aid future outage operations and resolutions. As overviewed in this section, MSOs make many customer-impacting decisions in real-time during CPOs and having an AI/ML framework to help direct these decisions could be beneficial. Section 3 will discuss more on the data storage and ML/AI engagements at Comcast.

 

## 3. Machine Learning on Historical Outages

The ability to leverage predictive models trained on historical CPOs can lead to more efficient resolutions of future CPOs. The various steps in the outage life cycle require different levels of validation and workforce planning based on dynamically changing data as the outage progresses. It is in these cases that machine learning can provide more confidence in validating different scenarios as it continues to learn and evolve from previous outages. As discussed in Cruickshank (2021), MSOs often have high fidelity data around CPOs and thus there is great potential to apply AI/ML techniques.

This section discusses what data related to each outage is stored and how that enables Comcast to run models that both allow for more effective action during subsequent outages as well as shows what steps can be taken prior to future outages that would minimize their impact on customers.

### 3.1. Creating a Historical Outage Data Base

When an outage concludes, a data structure is stored that contains information necessary to reconstruct the event, so that AI/ML models can be subsequently trained to learn and predict aspects of future outages. In what follows, we refer to this data structure as an 'outage object,' distinct from an 'outage event,' which is the clustering of offline customers at a given time.

At its highest level, an outage object stores summary information: the classification of the outage (CPO or HFC), the starting and closing time of the outage; the physical boundary enclosing all locations that experienced an offline event during the outage; the list of nodes and power supplies impacted; and the number of customers who experienced a service disruption. This information is useful for characterizing and comparing the overall impact of an outage. Figure 3 shows a snapshot of outages from a geospatial point of view for context about the meta data stored.

The outage classification is determined based on several factors. For example, in Figure 3, Outage A would be classified as a CPO due to the involvement of PSs. If Outage B happened at a similar time to Outage A, then it could be classified as a CPO based on proximity to Outage A, even though it has no PSs directly contributing. Finally, if Outage C happened in isolation at a separate time, it would be hard to classify it as a CPO without more context. A potential method for having a confidence factor when classifying CPOs with no PSs or nearby CPOs using historical data is mentioned in Section **Error! Reference source not found.**. The outages can evolve over time as the situation changes. For example, power providers bringing power back block-by-block. In these cases, an outage object should contain all the various states of a given outage, from start to resolution.

The next level of an outage object is a time series of outage events. For each observed timestamp and outage identifier, the outage object contains a list of the nodes and power supplies impacted; the number of customers offline; and the geographic boundary of the event. The boundary is derived from the locations of the offline customers and power supplies. This data can be used for further modeling and analysis efforts after the fact.

The final level of an outage object contains records for the components of the outages indicating the start and end time of their outage; their location (latitude/longitude); the node they are serviced by or power supply; and their status (offline for customers; offline or on standby for power supplies). Each record also stores the outage identification(s) the record is associated with.

 


**Figure 3 - Outage geospatial example**

More detailed information is available from the power supplies due to their continued telemetry, sampled every minute even while operating on standby. This includes detailed information about the power supply itself, such as the quantity of batteries and their make/model; the temperature; and the output voltage and power provided to the active devices in the OSP. Comcast also stores the individual battery voltages and the voltages over each string of batteries, allowing for critical insights into the operation and health of the batteries during their discharge. The PS telemetry that is stored for these historical outages are then used to train ML models that can predict battery run times during future outages as shown in Section 3.2

### 3.2. Power Supply Battery Run Time Modeling

By harnessing the rich battery run time telemetry from historical CPOs and leveraging the extracted historical performance data for each power supply, an advanced battery run time prediction model has been put into production. The model combines real-time telemetry with the aforementioned historical performance features to predict how long each individual power supply's batteries may be able to power the HFC plant until the batteries deplete and reach end of discharge (EOD). The current model is a Random Forest (RF) Regression model that was chosen as RFs are known to be capable of capturing complex relationships to outliers resulting in higher accuracy. Moreover, this model offers valuable insights through feature importance explanations and provides a confidence score in the prediction, making it useful in the field. This concept was first introduced in Ohnmacht and Stehman (2022) and the model accuracy has since improved due to the accumulation of more training data and enhancements to the features. The model in use is trained, validated, and tested on ~40K outages spanning over 10 months of data, resulting in a model that is informed by a wide distribution of possible run time events, empowering it with more predictive ability.

Figure 4a shows the current model performance against the first iteration of the model and a state-of-the-art physics-based model, described in Lin and Nispel (2022), on the test set of ~6k events, excluding events from the training set. It is evident that the new model performs significantly better overall but also on shorter run times, centering the raw prediction error closer to zero. Given the model's knowledge of historical events, it is outperforming the preceding model by 15% and the baseline physics-based model by 50%. The latest model was specifically tuned to make more accurate predictions for shorter run time situations. This accuracy improvement is seen in Figure 4b as the ML 2023 model is able to capture the trend of the shorter run times while previous models cannot. This was accomplished by adding a new set of features to the model that indicate if the power supply's batteries are not in a nominal state (many recent discharges, loss of commercial power while charging, etc.). This addition allows the model to make predictions that consider the current state of the batteries at run time.

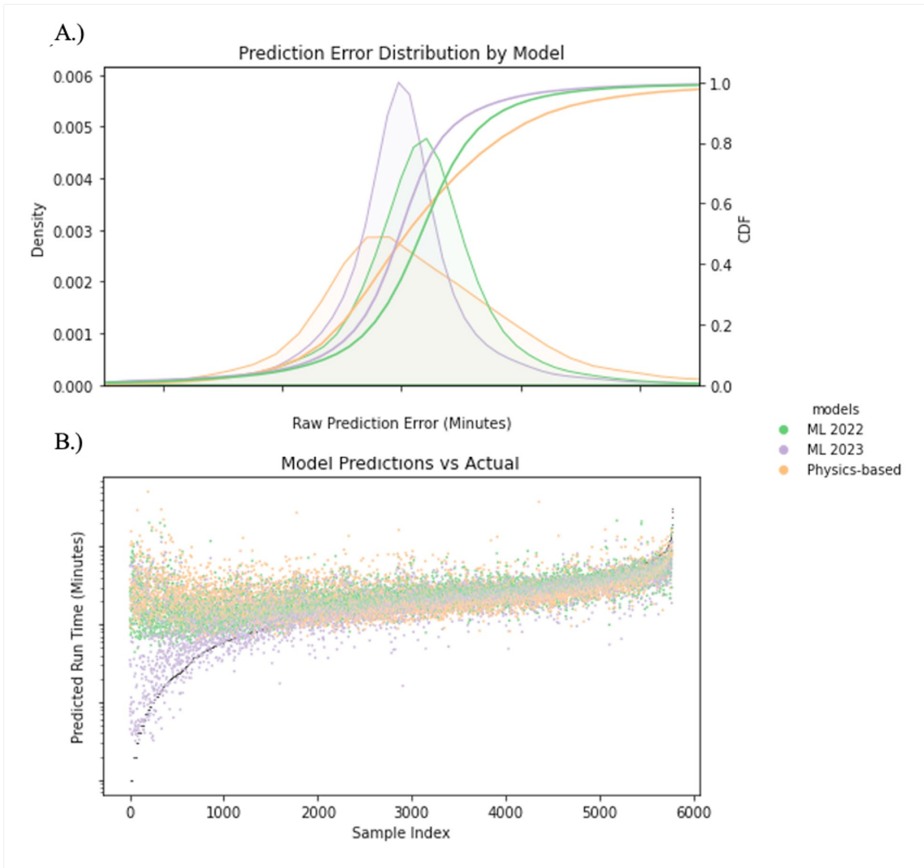


**Figure 4 - Battery run time model performance: A.) Distributions of model prediction errors B.) Raw predictions of each model (colored dots) compared to the actual run times (black dots)**

As with any production ML model, it is vital to have a model maintenance plan to avoid any sort of input data or model performance drift. From the most recent model, it is apparent that adding more events can improve accuracy. The current plan is to re-train this model on a quarterly basis, incorporating the most up-to-date distribution of possible outage scenarios while also exploring new features and methodologies for model improvement. There is also an internal view to monitor model performance that can eventually be incorporated into a pipeline for automatic model re-training.

### 3.3. Customer Impact Modeling

Once a historical database of validated CPOs is stood up, AI/ML algorithms can be deployed to find correlations and make predictions about future outcomes for similar outages. It is well known that the HFC plant and commercial power grid don't always overlap nicely. This discord can lead to various scenarios with customer connection status not aligning with the HFC plant status. The path to resolution may vary in these different scenarios and being able to predict beforehand how the customer connectivity in each outage scenario will unfold poses a huge head start in a path to resolution. See Table 1 for a breakdown of these scenarios, the most probable actions and potential customer communications.

 

**Table 1 - Actions for various outage states**

| Commercial Power Status | HFC Status | Customer Status | MSO Action | Customer Communication |
|---|---|---|---|---|
| In Outage | Online | Online | Monitor outage progression. Plan generator deployment around PS run time model prediction | Notify of current outage and that they may experience service disruptions if outage last longer than run time of batteries |
| | | Offline | Monitor customer. Likely lost power and cannot connect to service | Notify of current outage and that services are available. Include estimated time to power restoration |
| | Offline | Offline | Deploy generator and bring plant online ASAP | Notify of current outage and expected time to repair |
| Nearby Outage | Online | Online | Monitor nearby outage progression | Notify of nearby outage and potential disruptions |

From Table 1 there are a few interesting scenarios relating to customer experience. The first case is where some customers remain online from a DOCSIS standpoint even though the PS supporting their branch of the HFC plant is running on battery, indicating a loss of commercial input power. In this scenario, either the customer is on a different commercial grid than the HFC plant, or they have an alternative power source. These customers may not even be aware that there is an outage in their area and that their cable data services may be impacted. Being able to use advanced techniques to identify these cases and message customers about an outage and potential interruptions can go a long way in customer satisfaction.

The other scenario is when the commercial grid directly overlaps with the HFC plant. In these outages, most customers will show as offline while the PS is running on battery. This can completely change both operations and customer communications. Firstly, if all customers on a node are offline, even when the power supply is running on battery, that indicates that the customers share the same commercial power grid as the HFC plant. In these cases, if customers lose power, they might not know that cable data services are still available to them even though they lost power. Thus, in these scenarios, messaging a customer to let them know that cable services are still available could completely change their outage experience. Furthermore, if the AI model indicates areas where this scenario is pervasive over time, a proactive notification can be sent to customers indicating that an addition of a backup battery to their cable modem will keep them connected even during CPOs.

Comcast is developing these AI systems to learn trends about our HFC and CPO relationships using all the data and tools mentioned above. Figure 5 shows a plant topology view from a small example where correlation scores are given to customer locations on whether they share the same power grid as the HFC plant. In this example an entire busleg remained online while the PS was running on battery and went offline when the PS batteries depleted. Having a model that can learn from these scenarios and make predictions about the various states of customer connectivity during CPOs will significantly aid in outage resolution going forward.

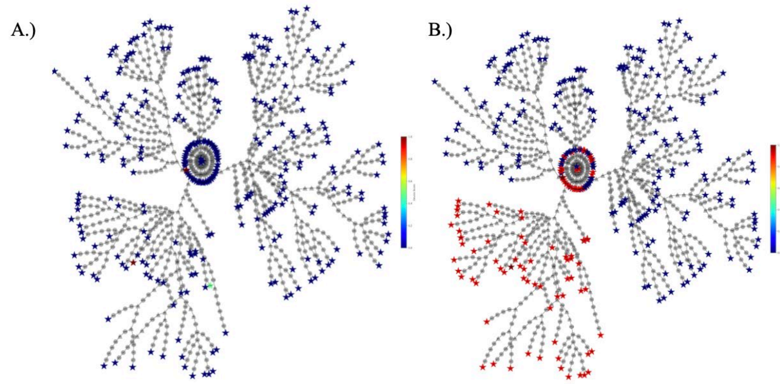


**Figure 5 - Customer correlation to power supply status: A.) customers online while PS is running on battery B.) customers offline when PS batteries deplete.**

## 4. CPO Dynamics and Long Term Planning

In this section, Comcast's approach to rethinking the CPO life cycle will be presented, highlighting the integration of AI and ML models. These models not only assist in resolving current CPOs but also play a crucial role in long-range planning to mitigate future outages. First in Section **Error! Reference source not found.**, an example CPO timeline is presented and discussed to showcase the power of building predictive models from historical outages. Then Section 4.2 will discuss the operational improvements that can be made with this type of approach.

### 4.1. Learning Patterns from CPOs

There are certain patterns that emerge when analyzing CPOs and understanding how to analyze these patterns can go a long way in being able to better action around future CPOs. Figure 6 shows a representative illustration of a typical CPO. This is a time series view of the typical components involved in a CPO i.e., counts of subscribers offline; power supplies running on battery; and power supplies whose batteries have reached EOD.

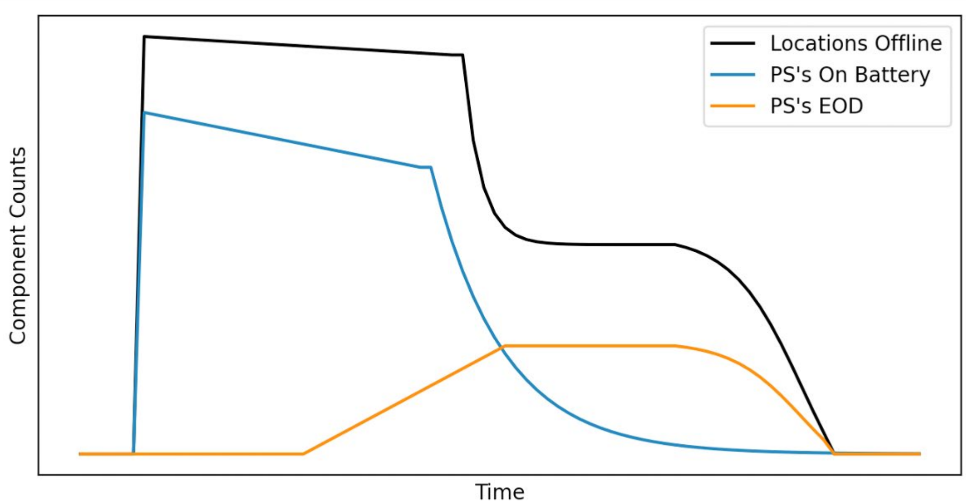


**Figure 6 - Timeline of a CPO**

 

There are a few features to note in this figure that will be relevant to modeling CPOs. The first feature being the initial impulse where subscribers and PSs lose commercial power almost instantly. As the power supplies run on battery, there tends to be a linear increase in the number of power supplies reaching EOD vs time. The number of PS's and customers that regain commercial power increases with time as the outage begins to resolve. It is easy to see that the offline location counts closely resembles the sum of the time series of the PS's on Battery and PS's EOD curves. This highlights the fact that some customer locations share the same commercial power as the power supplies and some locations have a different power source than the power supplies. For this reason, being able to model the relationship between customers, HFC plant and commercial power grid would allow MSOs to be more targeted and efficient when managing future CPOs.

Some important aspects worth modeling from this outage will be discussed as they can lead to increased confidence in identifying a potential future outage in the same area of service. Modeling which customers lose service at the same time will give an indication of which customer's share the same commercial power grid. In a similar fashion, knowing which customers regain service at the same time can help understand trends in resolving outages. Correlating customers to power supply status's is also of interest. Keeping track of which customers still have service when their power supply is running on battery can indicate that they have an alternative source of power. Analyzing which customers loose service when a PS runs to EOD can help inform power to HFC plant mappings and help in impact assessments of future EOD events, as mentioned in Section 2.3. Finally monitoring how long the Power Supply batteries last before reaching EOD can help in for future predictions of run time for subsequent outages, as discussed in Section 3.2.

In the event of a subsequent outage in the same area of service having models around the data discussed in this section can aid significantly in identifying and resolving the outage in a CPO standpoint. Understanding which customers tend to lose service together during CPOs will help add a level of confidence when declaring a future CPO. Having a model of which customers tend to regain service together can help when declaring an CPO closed and even indicate that the CPO is closed but an HFC outage remains. And finally having a baseline for how long a PS's battery lasted in a previous outage will be helpful for predicting its runtime in a future outage.

### 4.2. AI Enhanced CPO Operations and Planning

The previous section introduced several opportunities for AI/ML systems to add confidence in CPO identification and isolation as well as to help drive efficient field operations during outages. Figure 7**Error! Reference source not found.** shows a diagram of the overarching effort to have AI platforms not only impact real-time operations but also drive long-range planning.

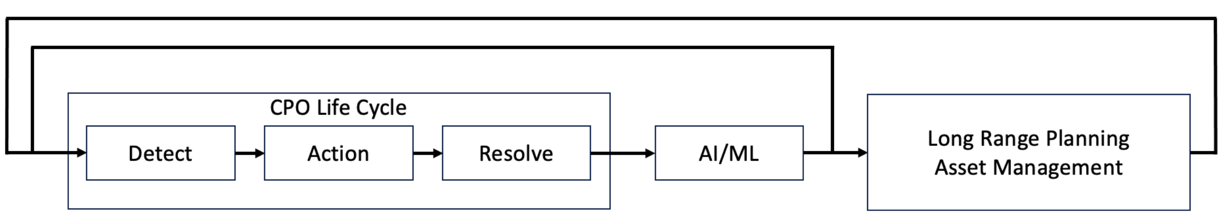


**Figure 7 - AI enhanced CPO resolution and OSP power asset management**

As showcased earlier in this section, while each CPO has its own features, CPOs don't tend to be random, and thus building models from historical CPOs can be useful for long-term planning. A few AI/ML initiatives that Comcast believes to be significant enhancements to the current state of the art are discussed in the remainder of this section.

From an operations standpoint, making targeted and calculated field decisions during CPOs is needed to mitigate customer offline time. Predictive models can aid in identifying CPOs and further distinguish between CPO and HFC outages. By modeling which customers typically go offline together in CPOs, a confidence score can be derived to help with the validation stage of future CPO lifecycles. Also understanding which customers share commercial power with the HFC plant allows for accurate proactive messaging and better planning in real time. This can be achieved by modeling which customers are typically online while being located inside an outage footprint.

With a rich historical database of CPOs, commercial power grid reliability and modeling could be leveraged for long-term planning. Building predictive models for commercial power grid performance and reliability can aid long-range planning on several fronts including hardening OSP power supply architectures and OSP asset placement. In the first use case, being able to model the commercial power grid from a reliability and geographic footprint standpoint would allow MSOs to set location-specific run time requirements for different regions of the HFC plant. The current PS backup battery systems may be over designed in some cases and under designed in others. Tailoring the PS run time designs to each location could help optimize OSP power supply efficiency. This information can then be made available to consumers through partnerships with CableLabs and platforms like Grid Metrics to continue improve customer experience across the cable and commercial power domains.

Another interesting use case is the potential to relocate OSP power assets to more reliable areas of the commercial power grid. This initiative has potential to offer savings on multiple fronts. Firstly, if the OSP power supplies can be moved to different and more reliable portions of the commercial power grid than the customers it serves, there is potential to have less HFC related outages and thus lower customer downtimes. From an energy usage standpoint, smart placement of OSP PSs can lead to less run time on battery, resulting in fewer re-charge cycles and further reducing the energy footprint of the OSP. Optimal placement of power supplies can also increase the service life of the batteries, resulting in a reduction in truck rolls and capital expenses.

## 5. Conclusion

This paper showcases how DOCSIS connectivity data and OSP PS telemetry can lead to advanced CPO detection, actioning, and resolution methodologies. Applying AI/ML models to a historical database of CPOs can lead to more efficient field operations during CPOs and aid in long-term planning efforts to further optimize the OSP power network. Comcast is currently running a ML PS battery run time model in production that significantly outperforms the state-of-the-art physics-based models. Comcast is also developing a novel CPO outage detection algorithm and creating a historical data base of CPOs for future ML/AI models. By applying novel AI/ML approaches in the commercial power domain, Comcast hopes to improve customer experience during CPOs by having more targeted, proactive customer communications and running a more efficient OSP power network, both physically and operationally.

# Abbreviations

| | |
|---|---|
| AI | Artificial intelligence |
| CPE | Customer premise equipment |
| CPO | Commercial power outage |
| DOCSIS | Data over cable service interface specifications |
| EOD | End of discharge |
| ETR | Estimated time to repair |
| HFC | Hybrid fiber coax |
| ML | Machine learning |
| MSO | Multiple systems operator |
| OSP | Outside plant |
| PS | Power supply |
| RF | Random forest |
| TSP | Telecommunications service priority |

# Bibliography & References


Ohnmacht, S., & Stehman, M. (2022). *Machine Learning and Telemetry Improves Outside Plant Power Resiliency for More Reliable Networks*. SCTE Cable-Tec Expo 2022, Philadelphia, Pa.

Cruickshank, R. (2021). *What's Smart About Smart Power?*. SCTE Cable-Tec Expo 2021, Virtual.

Lin, K., & Nispel, M. (2022). *A New Model for Plant Power and Health Estimation*. SCTE Cable-Tec Expo 2022, Philadelphia, Pa.

Peck, T., & Frankhouser, J. (2021). *Ensuring HFC Network Resiliency During Extended Utility Outages*. SCTE Cable-Tec Expo 2021, Virtual.