A Deep Learning Approach for Detecting RF Spectrum Impairments and Conducting Root Cause Analysis

A Technical Paper prepared for SCTE by

Kevin Dugan
Data Scientist
Comcast
1800 Arch Street, Philadelphia, PA 19103
719.493.2600
kevin_dugan2@cable.comcast.com

Justin Evans
CoreTech
Comcast
1800 Arch Street, Philadelphia, PA 19103
719.493.2600
justin_evans@comcast.com

Maher Harb
Director, Data Science
Comcast
1800 Arch Street, Philadelphia, PA 19103
215.990.8376
maher_harb@comcast.com

Table of Contents

List of Figures

List of Tables

1. Introduction

In proactive network maintenance (PNM), a goal of downstream spectrum analysis is to identify customer devices with impairments in RF (Radio Frequency) spectrum. Identifying the type of impairment is useful in determining the general causation, which is leveraged to infer the geolocation of the impairment's origin in the network. This paper describes the implementation of an automated, end-to-end solution that analyzes millions of sets of spectra and delivers PNM opportunities for the organization.

Utilizing wideband RF frequency response data from full-band captures (FBC), spectral impairments in the downstream are detected using a signature matching algorithm implemented as a 1-D convolutional neural network (CNN). The signature matching algorithm evaluates the set of spectrum data for customer devices on the cable plant. Certain impairments that originate in the cable plant, such as resonant peaks that can occur at an amplifier, impact the RF signal for multiple customers further downstream of the impacted component. One of the goals is to identify when multiple customers are experiencing the same signature, pointing to an issue in the network that can be resolved without individual visits to affected customers.

The impacted modems feed into a root cause analysis (RCA) algorithm that overlays the impaired devices onto a graph representation of the network topology. Through methods rooted in graph theory, the RCA algorithm narrows down the geolocation and system component(s) for the probable network device(s) where the issue originates. The impact of this workflow is an automated capability to identify not only customer-impacting issues, but also opportunities to proactively resolve issues before they become impacting across the entire network.

2. Background

2.1. Proactive Network Maintenance (PNM)

Utilizing full band capture data for the purpose of PNM initiatives to detect impairments in the downstream frequency response has been discussed in significant detail in multiple preceding works. Those contributions have directly influenced this automated system, by serving as a general roadmap for the planning, design, and implementation. The scope of this document starts with these works as the base knowledge and describes an automated system to enrich PNM opportunities [1-10].

2.2. FBC Data Characteristics

The basic unit of data for this system is the FBC of RF spectrum for customer devices where each spectrum sample consists of 8,704 values spanning 6 MHz to 1026 MHz. Spectrum features important to this work include downstream SC-QAM channels, guard bands, vacant spectrum, and pilots. Figure 1 illustrates a sample of a normal frequency response from a modem that maintains consistent power for occupied spectrum at the appropriate levels (approx. -17.1 dBmv).

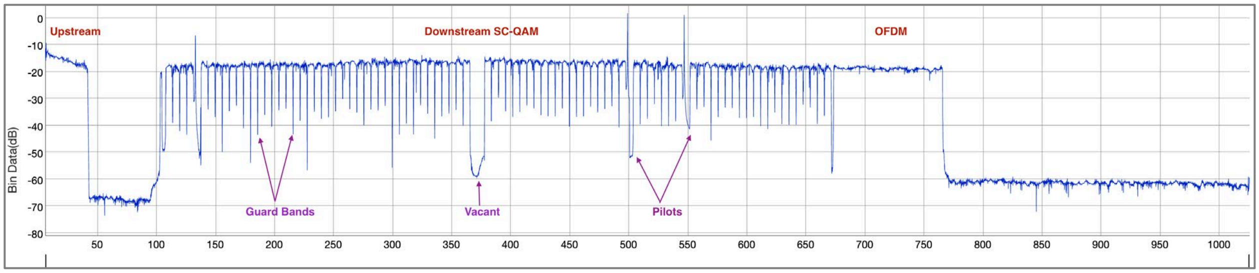

Figure 1 – An Example of a Normal Frequency Response

Downstream SC-QAM channels are defined as 6 MHz wide spans separated by guard bands. Along with pilots, guard band power values are not particularly useful in detecting impairments in many cases. Vacant spectrum represents a span where no services occupy one or more channels and can be identified programmatically by the power values and channelizing the spectrum. By reducing the spectrum only to occupied channel signal power, the signatures become more amenable to machine learning techniques that detect anomalous spectra.

2.3. Downstream Wave Impairments

Certain impairments in downstream SC-QAM channels can be identified in the FBC. For the scope of this implementation, the focus is on four impairments – standing waves (amplitude ripples), water in the cable, resonant peaks, and suck-outs. These patterns will be referred to as the ‘wave’ patterns. Figure 2 shows examples of each type of wave impairment. Within each plot, the solid green line at -17.1 dBmv and solid red line (bottom line) at -33 dBmv represent an appropriate range for frequency response values of occupied spectrum.

Standing waves, caused by impedance mismatches, are periodic in nature and generally extend across the entire downstream spectrum (Fox, et al., 2021). When water is introduced to the cable, an aperiodic wave is produced due to random attenuation and may also be associated with a negative tilt (Fox, et al., 2021). Resonant peaks are significant, narrow spikes in the spectrum caused by any number of reasons (i.e., cold solder joints or loose modules) on network devices (Cable Television Laboratories, Inc., 2016). A suck-out is represented as “a concave notch with sinusoidal boundaries with attenuation in amplitude/power caused by impedance mismatches evenly distributed through the network” (Cable Television Laboratories, Inc., 2016).

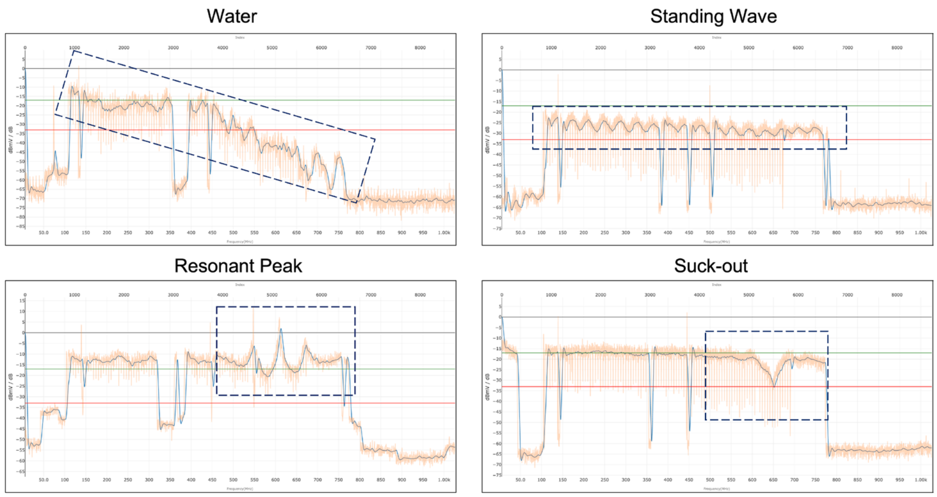

Figure 2 – Examples of Wave Impairment Frequency Responses

2.4. Network Topology as a Graph

A convenient way to analyze network topology is through graph theory. Comcast has mapped the "access network from the CMTS to the customer premise equipment (CPE)" into a graph structure "while incorporating all of the physical and logical elements that form part of the network" (Harb, Subramanya, Narayanaswamy, Walavalkar, & Rice, 2021). The graph facilitates the application of algorithms, such as lowest common ancestor (LCA), to network elements and their attributes. For example, clustering RF impairments on the graph gives the organization comprehensive knowledge about the network elements involved and a view of the common experiences amongst multiple customers. In turn, this knowledge is leveraged to deploy the correct resources to a specific physical location for resolution. In combination with network monitoring tools, an opportunity arises to automate the workflow from the current manual process.

3. Model Architecture

The adopted neural network for classifying RF impairments is a four-layer CNN that makes binary classifications for each of the wave patterns. Figure 3 represents the architecture diagram for building the pattern detection data model. Each of the 1-D convolutional layers uses a kernel size of five and a rectified linear unit (ReLU) activation. The convolutional layers are followed by a down sampling operation via max pooling. The max pooling operation calculates the maximum value in each section of the feature maps, pointing to the most present features. The increasing number of filters, as the CNN grows in depth, is attributed to the larger number of pattern combinations in each subsequent layer and using an increased number of filters allows the capture of more abstractions from the signal data.

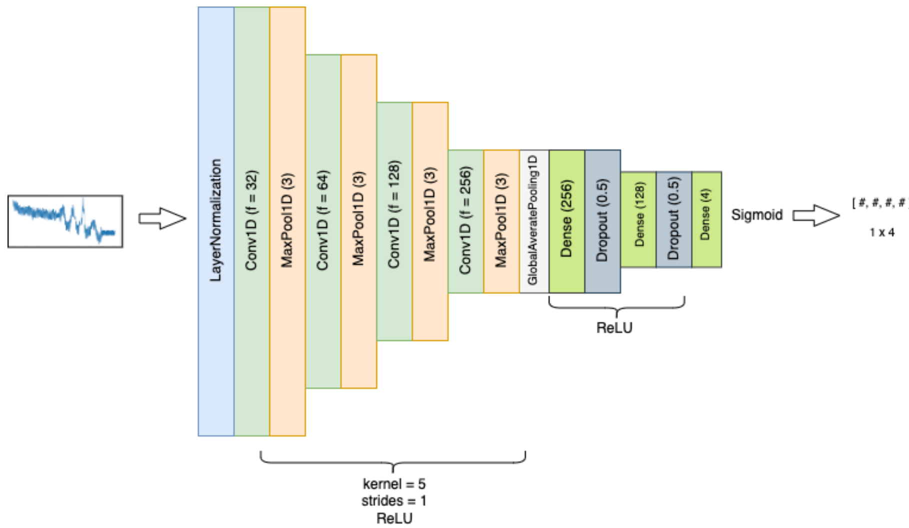

Figure 3 – The Neural Network Architecture

Multiple impairments may manifest on a single capture of spectrum. For example, for water to get into the cable, it needs an access point, such as a nick/chew in the cable, which would commonly cause standing waves (Fox, et al., 2021). Therefore, in certain cases it is difficult to differentiate among them, and their classification may impact the type of maintenance that would be referred to resolve the issue. It is more common to observe water in the drop cable or tap near the home; however, it may be observed in components for a larger group of customers and thus may be an issue that requires a network technician (Fox, et al., 2021). Being able to differentiate the impairments guides downstream decision-making processes.

Instead of structuring the model to calculate a single classification per spectra, the model predicts a probability for each type of wave. Taking this approach gives additional opportunities for analyzing the signatures of wave types within any given group and gives insights that improve the root cause analysis. Table 1 demonstrates possible classifications of multiple impairments.

Table 1 - Example Binary Classification

Sample Identifier	Water	Standing Wave	Resonant Peak	Suck-out
A	T	F	F	F
B	T	T	F	F
C	F	F	F	T
D	F	F	T	T

4. Model Training

The training dataset consists of 3,170 samples of labeled spectra from a population of 10,000 sets of spectra. The validation dataset consists of 500 hold-out samples. The labels were manually entered by a group of subject matter experts over the course of four weeks. The full collection of impaired and non-impaired spectra was acquired from an existing threshold-based detection algorithm that served as the basis for addressing this problem with machine learning methods.

Due partly to the limited training dataset and available resources, applying transformations on the data to reduce the noisy data points allows the CNN to learn appropriate classifications under the constraints. A downside of the transformations is the extended processing time for the feature extraction portion of the pipeline, processing millions of spectra for each iteration in production. As part of the model re-training cycle and growing the training dataset from validated classifications, we are optimistic that the preprocessing steps will be reduced in the future.

4.1. Labeling

To derive accurately labeled data, SMEs were presented with a random selection of both impaired and non-impaired sets of spectra via a user-interface (UI) adapted from a CableLabs initiative. Each spectra sample could be assigned any number of 13 labels. Figure 4 is a screenshot of the UI with a sample containing a standing wave. The two horizontal lines at -17.1 dBmv and -33 dBmv serve as visual indicators of the tolerable range in DOCSIS protocol to facilitate user interpretation.

At least two different users would be presented the same frequency response, with the usable training data samples meeting the criteria that more than one user assigned the same label. While this approach expedites much of the label validation effort, the overall number of samples is reduced because of requiring multiple labels from different users. Based on initial modeling experimentation, it was discovered that high-quality labels would be more useful than a few thousand additional labels whose quality was not checked as stringently. A follow-on experiment determined that by adding samples with a single label degraded the model’s classification capability.

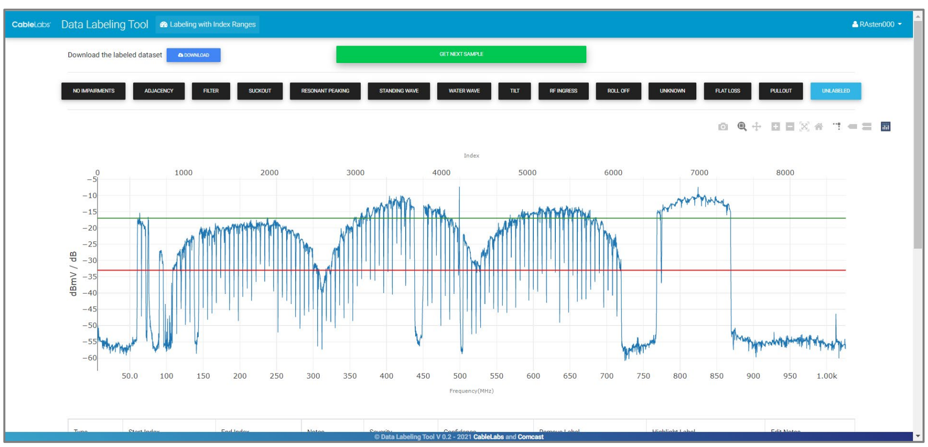

Figure 4 – The Labeling User Interface for Assigning Impairments

4.2. Transformations

The transformations applied to the raw data fall into one of two categories – data reductions and signal smoothing. Reduction logic reduces the number of data points per spectrum, while signal smoothing methods reduce the noisy parts of the spectrum’s signal including guard bands, pilot signals, and vacant spectrum.

4.2.1. Full Spectrum Smoothing

An initial smoothing algorithm is applied to the spectrum based on the channelization of power values. The static nature of downstream SC-QAM channel characteristics enables a smoothing algorithm that considers the signal between guard bands but leaves the guard bands intact. Since the data source for the smoothed channel data is from a visual tool, it maintained spectral characteristics for presentation in a UI.

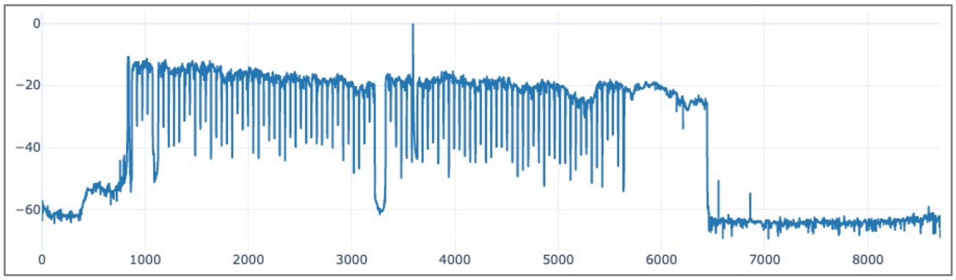

Figure 5 – Full Spectrum Smoothing

4.2.2. Truncate

The impairments under investigation all primarily reside in the downstream portion of spectrum. As a standardized transformation for all spectra, the usable spectrum for model training is defined as the span from 113 MHz to 748 MHz and may include OFDM spectrum values. As this applies to all spectra straightforwardly, any downstream services above 748 MHz are not considered for the wider bandwidth devices. Each sample in the transformed dataset now contains 5,420 values of smoothed spectrum.

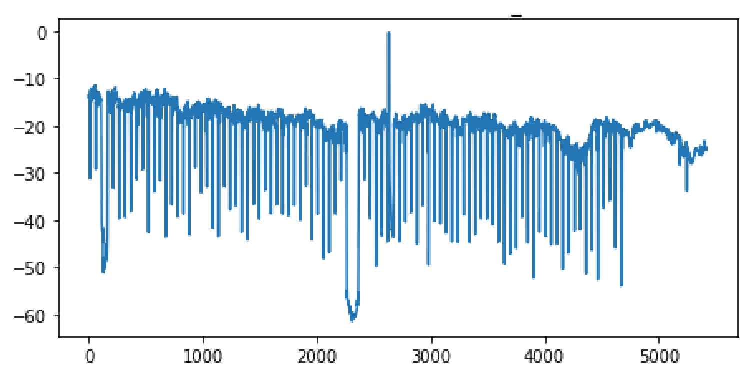

Figure 6 – Spectrum Truncated to Downstream

4.2.3. Binning

As a further reduction of data volume, the samples of 5,420 values each are placed into 2,000 evenly spaced bins. For each bin, the mean of the values inside becomes the updated spectrum values. This not only reduces volume, but also irons out excess data points that ultimately can be supplanted without degradation to the primary patterns in the spectrum.

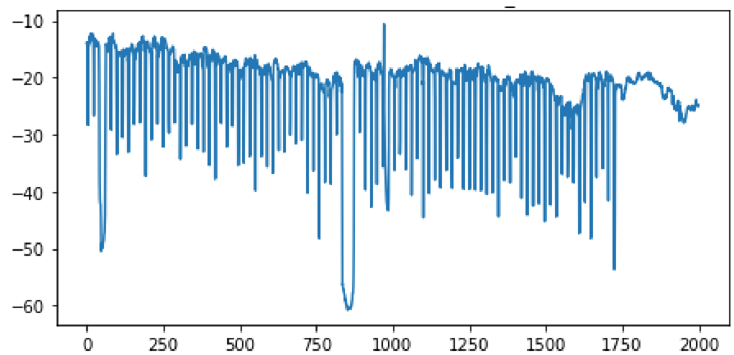

Figure 7 – Preprocessed Spectrum Binned to 2,000 Values

4.2.4. Guard Band and Pilot Removal

In early modeling experiments with labeled data, the data model would misinterpret guard band and pilot artifacts as characteristics of impairment types in many cases. Since this created a significant number of false positive results, it became necessary to smooth the values at these locations along the spectrum.

To remove the artifacts, the values are replaced with averaged values of the power immediately around the artifact – essentially creating a short linear regression line. A second approach to this for guard bands specifically would be using the channelized characteristics to replace the spectrum values. Both have proven to work well, and both achieve the desired transformational outcomes.

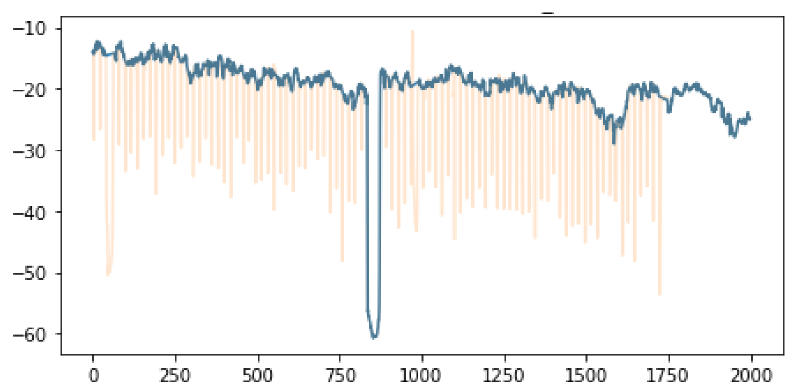

Figure 8 – Preprocessed Spectrum with Guard Bands and Pilots Removed

4.2.5. Vacancy Removal

The final transformation removes the spans within spectrum considered to be vacant – meaning there are no services in that span for any number of reasons. Vacancies are programmatically detected using thresholding techniques with the spectrum values being updated similarly to guard band and pilot artifacts. A linear regression is calculated between where the vacancy starts and ends; this matches the general trend of the spectrum at that location.

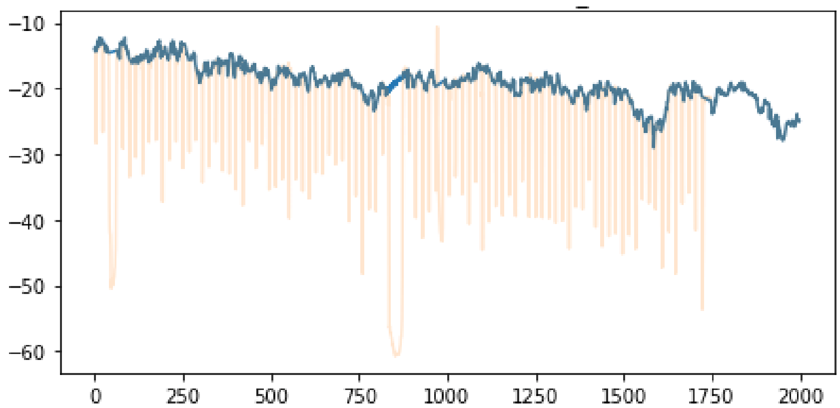

Figure 9 – Preprocessed Spectrum with Vacancies Removed

For very wide vacant spectrum, this approach may introduce a signature that resembles characteristics of one or more impairments. While the false positive rate is very low for these situations, discovering improved vacancy removal algorithms remains an important task. Again, the channelized spectrum data provides another strategy to identify vacant spectrum with high confidence and is the next step in this evolution.

4.2.6. Final Data Form

The final data form consists of 2,000 values, with noise removed for the model while retaining the signature's primary characteristics. From a scaling perspective, the transformations facilitate faster model training and inference operations, but at the expense of requiring more processing resources prior to involving the model.

The training data was lastly augmented by simply reversing the order of each sample, doubling the number of samples that retain the same pattern signatures but in different locations. Prior to training on the CNN, the data were passed through a normalization operation to improve consistency.

5. Model Performance

With a limited set of data to train the model, a 5-fold cross-validation architecture is implemented for estimating its generalization capabilities. The top performing cross-validation model is selected based on validation and training loss results, and then re-trained on the complete set of training data.

The model's ROC-AUC curve is 90% or greater for both the test data and for the hold-out validation dataset. The most confident classifications are the resonant peak and water signatures. Standing waves are the most difficult impairments to classify as certain signatures are similar to water or resonant peak signatures in some cases. Suck-out misclassification occurs primarily when artifacts remain from the vacancy removal process in which the leading or trailing edge was not properly cleared.

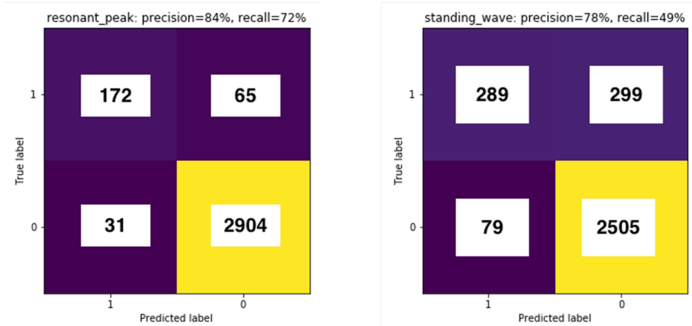

Figure 10 – Confusion Matrix for Resonant Peaks and Standing Waves

To calculate model lift, a dummy model (all samples labeled as no impairment) was evaluated in which the accuracy ranged from 81% for standing waves to 93% for resonant peaks. The impairment detection model outperformed the dummy model and created significant lift for each impairment in consideration.

Table 2 – Model Performance Metrics

Impairment	Training / Test AUC	F1 Score	Dummy Model Acc	Lift	Validation AUC
Resonant Peak	98%	0.78	93%	12.1	98%
Standing Wave	90%	0.60	81%	4.1	91%
Water	97%	0.74	89%	7.6	97%
Suck-out	94%	0.70	87%	6.7	91%

5.1. Model Predictions

Figure 11 highlights an example from each impairment with a correct prediction (TP), a false positive prediction (FP), and a missed prediction (FN) – reading column-wise. The red text in the middle row indicates the label applied to the sample during the labeling process described above.

The model performs exceptionally when the impairment type is straight-forward and mostly follows the definitions of the impairment. The model is less-confident when a frequency response indicates the presence of complex waves – that is, when multiple impairments or other conditions show characteristics of multiple anomalies. An example is the FP sample for standing waves shown in Figure 11. The frequency response has strong characteristics of a standing wave and weaker characteristics of a water wave, giving rise to a complex wave type. Since water waves are akin to standing waves, an assessment of this spectra is that there is moisture present in a standing wave, but not enough to induce drastic random attenuation. This determination impacts the root cause analysis algorithm, explained later in this document.

Another factor influencing model behavior is the volume of data and the quality of labels. The limited dataset lacks signature diversity, making it difficult for the model to interpret frequency responses that are considered complex. The labeled data also contains contradictory labels for similar-looking samples, demonstrating that interpretations of spectra can differ among individuals.

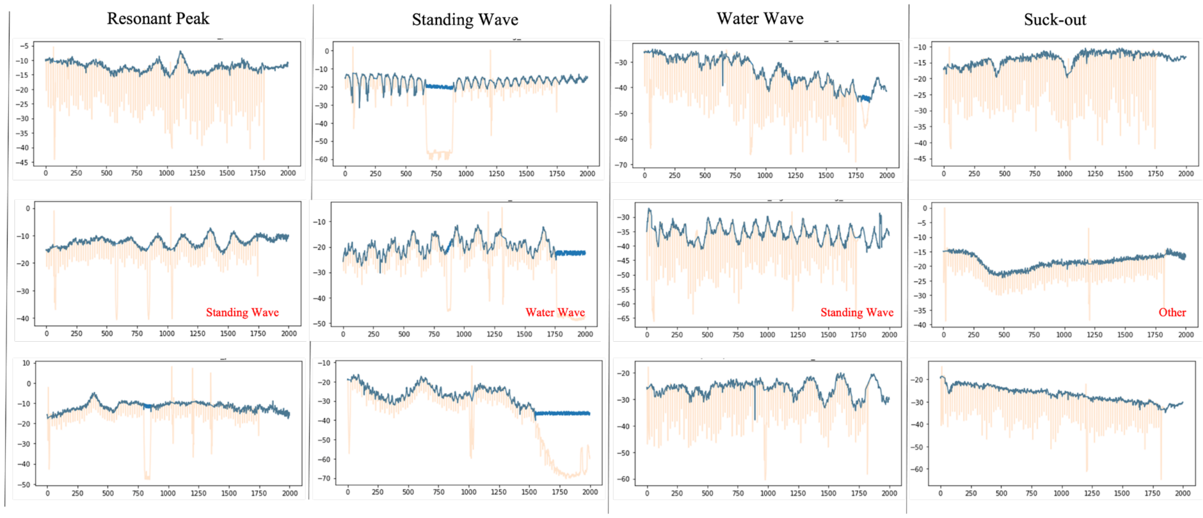

Figure 11 – Classification Examples (True Positives, False Positive, False Negatives)

5.2. Model Improvement

To obtain a larger training dataset for model recalibration moving forward, a system was established that collects random samples of predictions from each pipeline iteration, feeds them into a UI where a user validates the prediction, and the sample gets added to the training data. Figure 12 represents the workflow for retraining the model with additional samples.

The diversity in the initial training dataset was limited due to the low volume of available samples. To correct the model for any shortcomings related to this, obtaining validated predictions is critical for growing the training data. Additionally, the validation system can be filtered to a particular impairment so that any imbalances in the training data may be addressed by validating more quality samples of specific impairments.

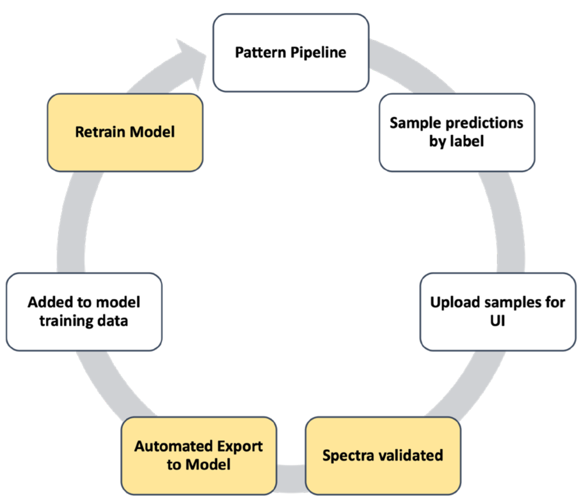

Figure 12 – Model Recalibration Cycle

In the validation process, a user can observe the original frequency response in addition to the data passed through the model. This has been helpful in catching edge cases with the preprocessing in which unwanted artifacts remained under certain conditions. In the vast preponderance of cases with an unexpected prediction, the artifact had a direct impact on the visual representation of the sample as it passed through the model.

6. Root Cause Analysis

An RCA algorithm was developed that leverages the graph representation of the network topology. Conceptually, the RCA examines a group of impaired modems within the context of seeking common network elements in the topology and calculates probabilities that each common network element may be the origin of the impairment. The algorithm surfaces results of the RCA to network monitoring tools, so that as PNM opportunities arise, the appropriate action can be taken to improve network conditions and reliability.

6.1. Methodology

The inputs into the algorithm consist of the following primary elements:

1. Type of impairment;
2. Impaired device list; and
3. Total device list in the grouping that reported frequency responses (a fiber node, for these purposes).

The graph of the topology for the node is extracted and a lowest common ancestor (LCA) search is performed for the impacted devices. The LCA makes different determinations on network elements depending on the type of impairment. Only amplifiers are considered when evaluating resonant peaks. For water in the cable, only customer drops are aggregated – meaning each LCA will point to a specific customer location and not a network element. Standing waves and suck-outs do not restrict which network elements are considered.

Standard F1 scores are calculated for each of the network components available in the topology as shown in Figure 13.

$$Precision = \frac{\#\ impaired\ devices\ on\ network\ element}{\#\ devices\ reporting\ spectra}$$

$$Recall = \frac{\#\ impaired\ devices\ on\ network\ element}{total\ \#\ impaired\ devices}$$

$$F1\ Score = \frac{(2 * Precision * Recall)}{(Precision + Recall}$$

Figure 13 – Standard F1 Score Calculation for Network Elements

The LCA scores are ranked in descending order, with the possibility that multiple vertices (network elements) have the same score. This occurs in situations in which a sequence of vertices has the same number of impaired devices and the same population of devices that reported spectra – i.e., when multiple networks elements are connected with no additional modems in between them. To break the tie, the element with the longest path back to the CMTS is selected, as that element is the one closest to the impacted devices.

6.2. Annotating the Plot

When plotting the data for visual inspection, both physical and logical network elements are represented using the color scheme shown in Figure 14 (left). The figure also describes the highlighting scheme that represents the status of individual customer devices.

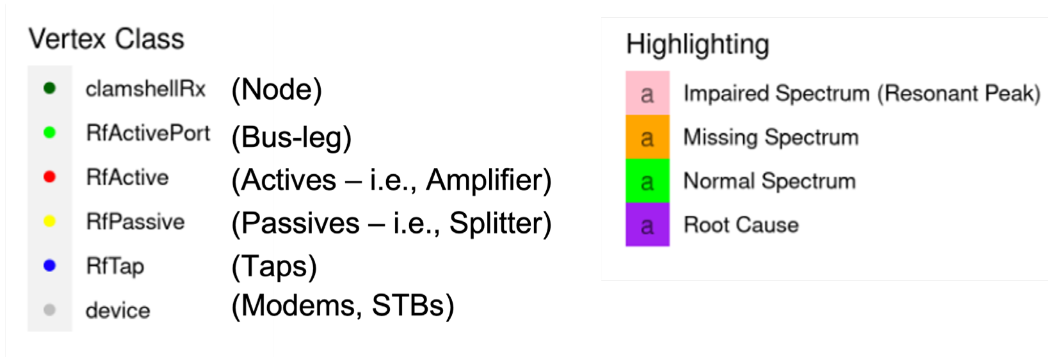

Figure 14 – Color Scheme Legend for RCA Graph Plots

The device highlighting provides visual clarification for where the impairments exist within the larger collection of devices available in the topology. Not all devices report spectrum, and since the LCA scores consider only the reporting devices, identifying them on the plot assists in validating the RCA results. High scoring vertices will contain few instances of normal spectrum and have more devices that are impaired or did not report spectra. The lowest common ancestor (highest scoring vertex) is highlighted in a unique color and has an annotation attached that describe the type of network element and the unique identifier.

6.3. RCA Results per Impairment

The initial problem statement for developing the RCA algorithm using the network topology was an issue commonly seen in amplifiers that causes resonant peaks to present in FBC data. For this reason, the RCA algorithm only calculates the score for each of the amplifiers found in the node's topology. In Figure 15, the modems impacted by resonant peaks are clustered on the left side of the graph (highlighted in pink), with no modem in the cluster reporting having normal spectrum. Therefore, the amplifier highlighted in the figure has an RCA score equal to the maximum possible.

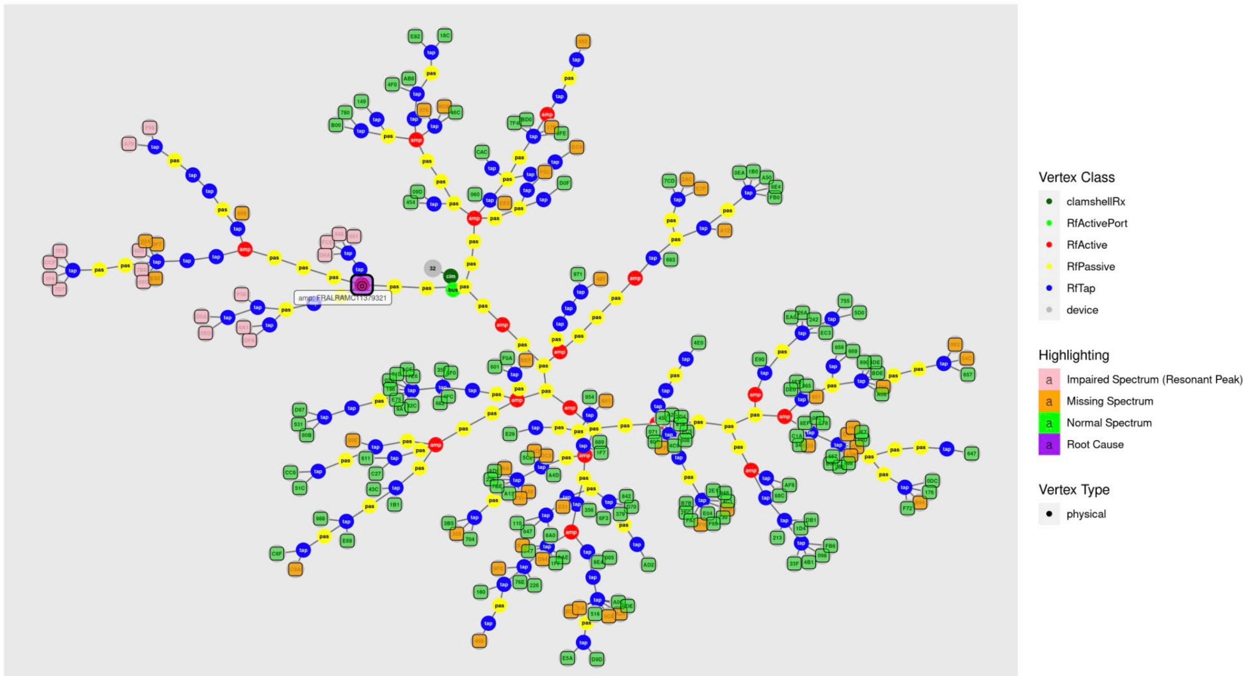

Figure 15 – Resonant Peak RCA Graph Plot

Standing waves are a result of impedance mismatches caused by damaged cables, breached cable jackets, improper connections, and animal chews, to name a few. Standing waves may be manifested either in the plant or in a customer's home. The RCA algorithm considered only the standing wave events that likely originate in the plant or at a multiple dwelling unit (MDU) by considering a minimum number of impacted devices. In Figure 16, the highest scoring vertex was a passive network device under which the preponderance of modems showed impairments alongside only a few with normal frequency responses.

The topology edge data includes properties about cable lengths between vertices. While not yet implemented in this system, the capability exists to calculate the length to the voltage reflection in the cable from a starting point. With the known length, it is then possible to further refine the LCA calculation by considering only common network elements with a minimum cable length of the known distance to the reflection in the line.

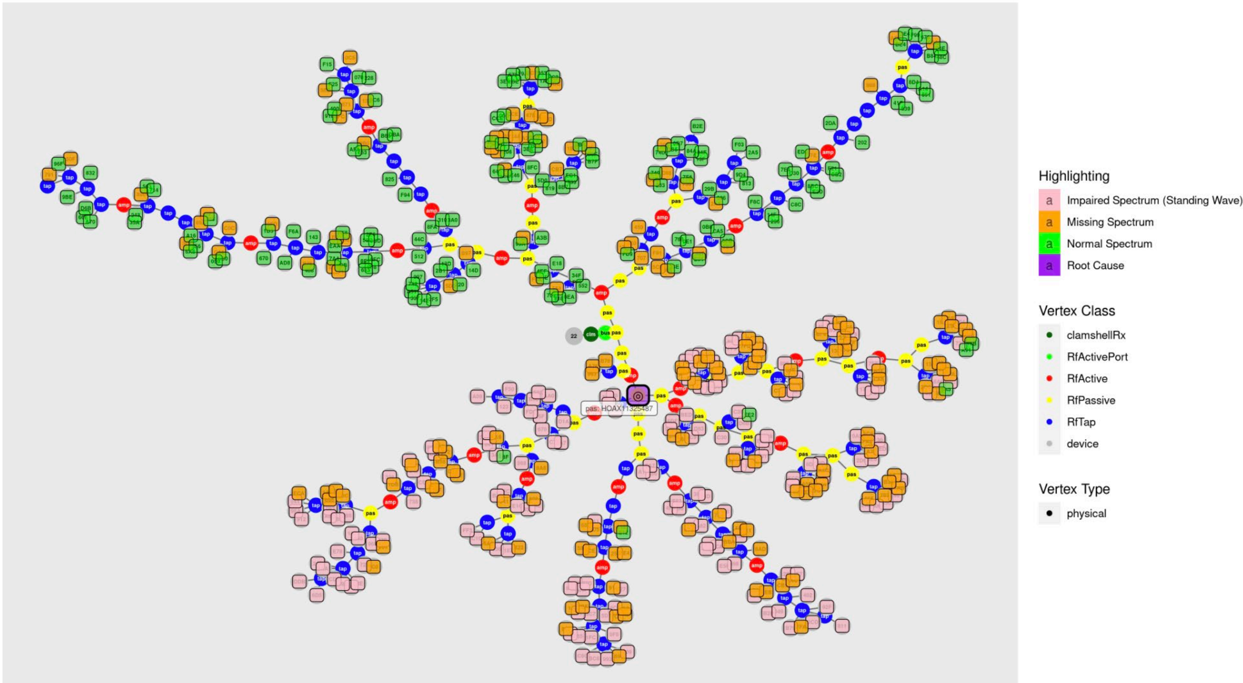

Figure 16 – Standing Wave RCA Graph Plot

Suck-outs are commonly caused by "mechanical or grounding issues in active or passive network elements such as seizures, connectors, lids, or fittings" (Cable Television Laboratories, Inc., 2022). Similar to the diagnosis of standing waves and resonant peaks, a minimum number of impacted modems were required to pass through the RCA. The identified vertex in Figure 17 was a tap under which some modems had normal frequency responses, and some showed the suck-out impairment.

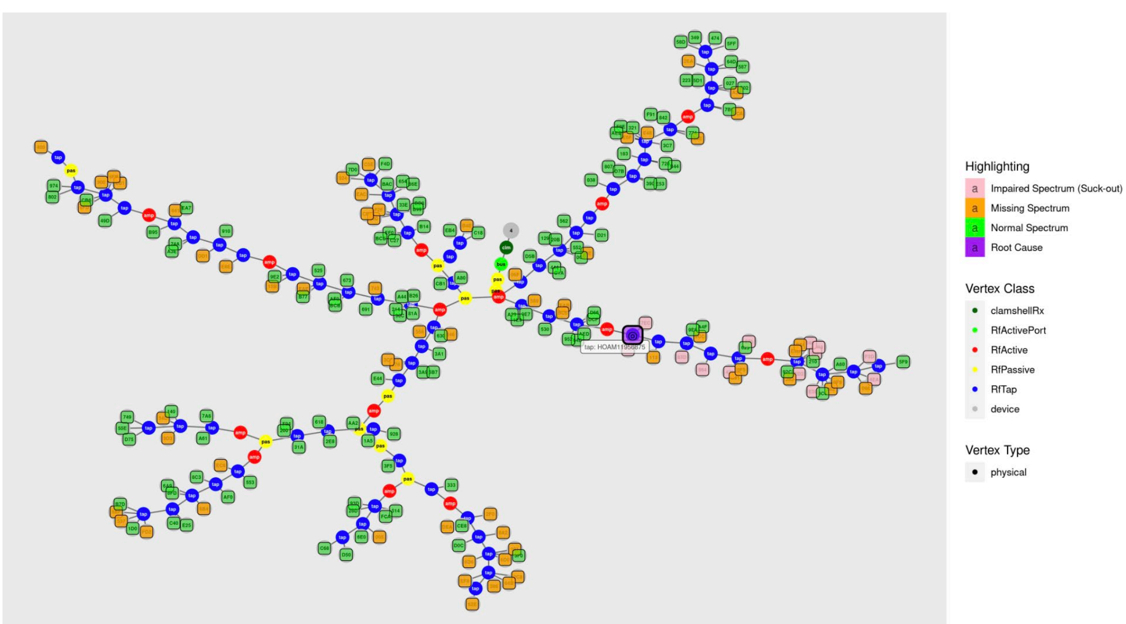

Figure 17 – Suck-out RCA Graph Plot

Water waves are most identified in customer drop cables and taps (Fox, et al., 2021). The RCA algorithm here only considers the drop locations to group the devices in the home/MDU into a single event. This was an intentional implementation decision while the clustering algorithms for water waves remained under development and will be used in future iterations.

7. Conclusion

Frequency responses from FBC data have been used extensively to identify certain impairments that impact network reliability. Additionally, multiple initiatives have surfaced deep learning models designed to automate detection of signal impairments in different types of signals. This implementation takes an additional step by automating the triangulation amongst network elements to narrow down the origin of the impairment through a network topology graph representation.

The deep learning model described here is but one form the model can take. In the development of this system, multiple models that produced quality results were evaluated. This detection model is an improvement over currently known threshold-based algorithms, often finding anomalous frequency responses where threshold-based algorithms miss them. The advantage to the previous algorithm is that it reliably finds moderate to severe cases; however, it is unable to differentiate among some wave types so they are categorized under one label. This impairment detection model not only differentiates the waves, but also has demonstrated the capability to detect both obvious and subtle impairment signatures.

By leveraging network topology as a graph, more advanced analytics are possible to understand the scope and impact of impairments on collections of modems. The lowest common ancestor algorithm is one of several forms of analysis that can be done. Through an understanding of RF impairments and their causes, a root cause analysis performed on network topology generates confident assessments of where within the plant the issue originates. The impact is better visibility into network conditions and improved support for the technicians that deploy to resolve issues in the network and at customer's homes.

Abbreviations

CNN	convolutional neural network
CPE	customer premise equipment
dBmv	decibels relative to one millivolt
FBC	full-band capture
FN	false negative
FP	false positive
LCA	lowest common ancestor
MDU	multiple dwelling unit
OFDM	orthogonal frequency-division multiplexing
PNM	proactive network maintenance
RCA	root cause analysis
ReLU	rectified linear unit
ROC-AUC	receiver operator characteristic – area under curve
RF	radio frequency
SC-QAM	single-carrier quadrature amplitude modulation
SME	subject matter expert
TP	true positive
UI	user interface

Appendix

Network Topology

Network topology can be accurately described as a directed acyclic graph or DAG. This allows for easier and more efficient algorithms to be used without the need to cycle check. A DAG is a special type of graph that allows traversal in one direction without the chance for a connection to "backtrack" to a previous vertex. The reason for this is plant topology does not loop back on itself and upstream edges will always direct to the devices closer to the CMTS in the plant.

The two primary plant topologies revolve around having an analogue CMTS doing all the work vs having a distributed access architecture where the load is balanced across the node. These two topologies do pose a challenge when creating the topology map as different devices have similar functionality yet different naming schemes. Nodes with vCMTS (virtual Cable Modem Termination Stations) have no amplifiers and scale incredibly well leading to easier changes in the plant.

Creating the graph that this work depends on took many teams a lot of effort in order to overcome the challenges required. In order to flesh out the graph, multiple data sources needed to be pulled from and synced. To achieve this the team used a property-graph architecture as well as Apache Tinker pop to perform the aggregation queries, look-ups ext. The graph database ROCI is still not fully mature but has proven extremely valuable in many projects including our own.

Bibliography

Cable Television Laboratories, Inc. (2016). *PNM Best Practices: HFC Networks (DOCSIS 3.0).*

Cable Television Laboratories, Inc. (2022). *Primer for PNM Best Practices in HFC Networks (DOCSIS 3.1).*

Ferreira, J., Harb, M., Subramanya, K., Santangelo, B., & Rice, D. (2020). Convolutional Neural Networks for Proactive Network Managment. *SCTE 2020 Fall Technical Forum.*

Fox, K., Rupe, J., Willimas, Tom, Zhu, J., Hranac, R., Zedan, N., . . . Wolcott, L. (2021). Water Can Run, But It Can't Hide. *SCTE 2021 Fall Technical Forum.*

Harb, M., Subramanya, K., Narayanaswamy, R., Walavalkar, S., & Rice, D. (2021). How Network Topology Impacts Rf Performance: A Study Powered By Graph Representation of the Access Network. *SCTE 2021 Fall Technical Forum.*

Hranac, R., Campbell, C., Fish, R., Kolze, T., Kristoffersen, E., Medlock, J., . . . Wolcott, L. (2020). Full Band Capture Revisited. *SCTE 2020 Fall Technical Forum.*

Virang, D., Chari, S., & Kraiman, S. (2021). Machine Learning for RF Impairment Classification. *SCTE 2021 Fall Technical Forum.*

Vishnyakova, A., Mahajanam, R., O'Dell, M., Merkle-Tan, M., Hay, C., & Pham, L. (2021). Right Technician at the Right Time. *SCTE 2021 Fall Technical Forum.*

Volpe, B., & Ottlik, B. (2021). Machine Learning and Proactive Network Maintenance: Transforming Today's Plant Operations. *SCTE 2021 Fall Technical Forum.*

Wolcott, L., O'Dell, M., Kuykendall, P., Gopal, V., Woodrich, J., & Pinckernell, N. (2018). A PNM System Using Artificial Intelligence, HFC Network Impairment, Atmospheric and Weather Data to Predict HFC Network Degradation and Avert Customer Impact. *SCTE 2018 Fall Technical Forum.*

Zhu, J., Sundaresan, K., & Rupe, J. (2020). Proactive Network Maintenance using Fast, Accurate Anomaly Localization and Classification on 1-D Data Series. Louisville: IEEE. Retrieved 2022, from https://arxiv.org/abs/2007.08752

