Right Technician at the Right Time
Using Machine Learning to Predict Network Maintenance Issues

A Technical Paper prepared for SCTE by

Anastasia Vishnyakova
Machine Learning Engineer
Comcast Cable
1800 Arch Street, Philadelphia PA 19103
267-260-3277
Anastasia_Vishnyakova@comcast.com

Rama Mahajanam
Director of Machine Learning
Comcast Cable
1800 Arch Street, Philadelphia PA 19103
267-260-4237
Rama_Mahajanam@comcast.com

Mike O'Dell
Director Network Maintenance
Comcast Cable
Virtual Location
412-417-0481
Michael_Odell@cable.comcast.com

May Merkle-Tan
Sr. Researcher, Machine Learning
Comcast Cable
4100 E Dry Creek Rd, Centennial, CO 80122
no contact number
Heng-RuMay_Tan@Comcast.com

Catherine Hay
Manager, Process Management
Comcast Cable
676 Island Pond Rd, Manchester, NH
603-222-7420
Catherine_Hay@cable.comcast.com

Lisa Pham
Machine Learning Engineer
Comcast Cable
1800 Arch Street, Philadelphia PA 19103
267-260-4319
Lisa_Pham@comcast.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

1. Introduction

When a customer encounters a service problem that cannot be resolved by regular triage with a call or chat agent, a fulfillment truck roll is scheduled. A standard technician, who specializes in repair issues located inside the premise and/or related to the service drop, performs onsite troubleshooting. If they discover an issue within the broader service network typically impacting multiple customers, they escalate the problem to a line technician. Line technicians expertise is in repairing and maintaining the Outside Plant (OSP) network, such as the nodes, amplifiers, passives, and hardline cables that provide service to multiple customers. This two-step troubleshooting practice incurs the cost of dispatching both types of technicians when an OSP network maintenance impairment is determined to be the root cause of a customer's service issue. More importantly, the follow-up escalation often delays the problem resolution for the customer.

In this paper, we will show that resolution efficiency can be improved by harnessing machine learning (ML) to predict when a customer's service issue has a high likelihood to be escalated to the line technician. Our approach leans on network and device telemetry as well as monitoring processes that assess the integrity of our service network (such as checking for outages, impairments, and performing problem segmentation). A key component of our ML modeling is the integration of graph features derived from network topology, that help identify issues related to equipment within the network that serves multiple customers. This model is in the trial stage and is being tested by selected regions. We also discuss how we can evaluate model precision to incur savings for implementing the model. Finally, we describe how the model will be integrated into the internal troubleshooting software.

2. Acknowledgements

We would like to acknowledge contributions of several Comcast colleagues. Michael Kreisel, who has since left Comcast, developed the initial RTM (Refer to Maintenance) model and the topology data. The extended Applied AI engineering team contributed to the setting up and maintaining of the data pipelines supporting the model. We would like to thank Kevin Bohinski, Ryan March, Likhitha Thebatni, Pat Dwyer, Nick Pinckernell, and Koundinya Venkata Sai Ravulapati for supporting the data structure for the model. We thank Applied AI researchers Yonatan Vaizman, Tianwen Chen, Fan Liu, Navdeep Jain, Scott Rome, Joshua Jackson, and Zhipeng Liu and director of Machine Learning Hongcheng Wang for their expertise and advice on the model development. We are grateful to our intern Vishnu Sharma who has supported model operations and contributed to the feature development. We thank Garret Beatty for his work during the Fall of Code program. We thank our product owners Dave Monnerat and Jason Stevens for being champions of the model among stakeholders. Background on plant operations and tools was obtained from interviews with many Comcast leaders and experts. We want to thank Larry Wolcott, James Sayer, Scott Johnston, Gary Ventriglia, Reid Downey, Scott Shrader, Andreas Lebaudy, Hari Palaiyanur, and Rich Aleong. We thank our field operations colleagues Ray Stewart, Mark Bosteder, Josh Halbrook, Nicole Atherholt, Jason Fletcher, Ed Reece, Brandon Yawn, Karen Washington, Robert Millis, Marcellis Price, Mike Mahaney, Andrew Herr, and Marc Tucker for their support during the model trial. We are grateful to Leslie Ellis, Larry Wolcott, and Jan Neumann for their edits and revisions.

3. Plant Operations

If recent global experiences have taught us anything, it is that connectivity to communications, information and entertainment are essential to modern life. There are many different technologies available to consumers for these needs, with cable TV / broadband networks being among the most common and widely available service delivery network types. During the pandemic of 2020, many service delivery providers, ourselves, and others, experienced significantly increased demand for data, as well as increased expectations for reliability and availability. Energy is still being invested in building scalable solutions to meet the increasing demand for bandwidth to the end user. For the purposes of this paper, we are going to examine and define a typical system, dividing the network into several segments, primarily the ISP (Inside Plant), OSP (Outside Plant), and Customer Premise.

The ISP, also commonly referred to as the Headend, is the facility where signals are generated, received, and processed from content providers, as well as from Customer Premise Equipment (CPE), like broadband gateways and set-tops. It is typically a climate-controlled facility where the equipment is housed to manage the routing of video, data, voice, and other signals to their intended destinations. The Headend also houses CMTS (Cable Modem Termination System) devices for broadband connectivity.

The actual customer premise – the household, multiple dwelling unit (MDU) or business -- can have slight variations in definition, primarily around where the demarcation point exists between the premise wiring and the Outside Plant. In some cases, the outside drop cable that connects the hardline cable tap to the premise is considered to be customer premise wiring, and in other instances, the customer premise wiring is defined as the customer side of the grounding or bonding block. In the context of the Refer to Maintenance process, this demarcation is an important distinction. For this paper, we will define the premise as any and all wiring from the hardline cable tap port to the Customer Premise equipment (cable modems, WiFi gateway, video set top boxes, etc.), including all coaxial cables, amplifiers and splitters in the customer premise wiring.

We can now define the Outside Plant as the transmission medium(s) transporting signals between the Headend or Hub and the hardline cable tap port. This would include the fiber and coaxial cables, as well as the active components (nodes and amplifiers) and passive components (splitters, couplers, and taps) between the Headend and the customer premise.

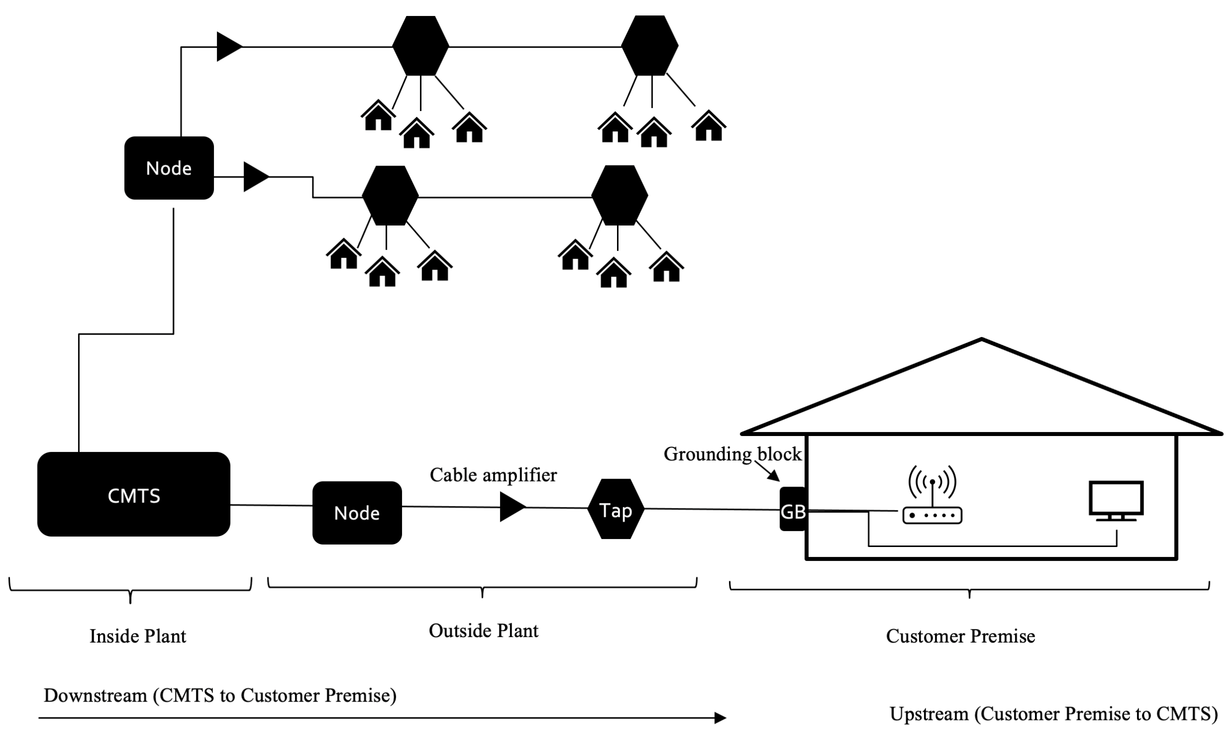

Figure 1 – Plant Diagram

As mentioned previously, this portion of the network is evolving rapidly to meet the increasing demand for bandwidth by the broadband consumers. The Hybrid Fiber Coax (HFC) portion of the network is becoming increasingly diverse as new technologies are developed and deployed, while previous generations of hardware are continually optimized to meet the signal quality and bandwidth needs of the customers. HFC networks built in the late 1990s can still reliably provide service to customers today. Even so, the components that make up the HFC section of the network have remained largely unchanged. To that end, we need to consider the failure points that are common to all generations of HFC networks.

The headend or hub locations that we mentioned have also undergone technology evolutions, but fundamentally, their purpose is to aggregate all the signals designated for a service area, and then modulate them into an optical format for transmission over a fiber optic transport network to a fiber optic node. The node translates that optical signal into an electrical signal and directs that signal into coaxial cable through a series of mechanical connections. The coaxial cable then transports those signals toward the target customer population through another series of mechanical connections via coaxial cable, connectors, couplers, splitters, and finally to the cable tap. The cable tap diverts a portion of that electrical signal to the customer service drop through another mechanical connection at the tap port. The service drop then transports the signal through yet another series of mechanical connections such as bonding or grounding blocks, splitters, amplifiers, connectors, and wall plates, until the intended signals arrive at a cable modem, set top box, or other customer premise equipment. In this example, we outlined many of the connection points that a signal must traverse in a single direction to reach the customer premise equipment, but this example did not include any active devices (amplifiers). This example of a single fiber node feeding a node service area without requiring any amplifiers is often referred to as a Node + Zero (or N+0) architecture, meaning there is a node, but no amplifiers to "boost" the signal to send it further into the network. Amplifiers allow operators to serve a larger area without adding more fiber cables or nodes. It is not uncommon for operators to still maintain HFC networks with N+6, or fiber optic nodes trailed by as many as 6 additional amplifiers, in cascade, required to reach all customers in the node service area.

Thus far we have only discussed the service delivery of signals to the customer in a single, or one-way direction. That is, signals are delivered from the headend or hub to all the customers in the node service area, or "broadcast". The reality, of course, is that there are signals flowing in both directions, from the headend to the premises, and conversely, from the premises back to the headend, aggregated via the HFC network components. However, all signals, whether upstream or downstream, can be adversely affected by impairments at any point in the HFC network. The degradation of the signal level or quality can be experienced very differently by the customers and our intelligence tools, depending on which component in the transmission medium is damaged or compromised.

The HFC network is a complex yet robust asset. It is arguably one of the most important assets for any operator, and as such, much effort and energy is invested in keeping it healthy and performing at an optimal level. As we described at a very high level, there are a very large number of components that must individually be uncompromised, and collectively maintained to ensure the physical, mechanical, and electrical integrity of the overall network.

Fortunately, the maintenance of these networks has gotten much more efficient over the years. While many of the troubleshooting practices have changed very little, the telemetry we are able to collect from CPE and other assets in the HFC infrastructure have allowed us to better identify those areas which are exhibiting symptoms of impairment. From that additional telemetry data, we can prioritize our response by customer impact potential, and dispatch our Line Maintenance teams to the right area to repair the network and remove the impairments. Previously, plant maintenance was either 100 percent preventative, meaning we would sweep or inspect the network on a scheduled basis, or it was reactive, meaning we responded to an area when our customers reported an issue. Using the data points that will be discussed later in this paper, we are able to collect metric data points from devices located in different points of the network and infer when those symptoms are caused by an impairment in the physical transmission medium.

4. Remote Network Telemetry Measurements

DOCSIS (Data Over Cable Service Interface Specification) technology is supported by comprehensive data collection systems that offer insights into the state of the access network remotely. Comcast has developed and operates several intelligence tools supported by DOCSIS specifications to help identify network impairments and prioritize repairs. Wolcott et al [1] wrote about the role of the network remote telemetry in the Proactive Network Maintenance (PNM) systems:

… [with] the capability of nearly all CPE [Customer Premise Equipment] to report some level of intelligence back to the tool sets, and [we] have a comprehensive view of all corners of the network. No longer is plant performance informed by a small quantity of DOCSIS channels on a relatively small number of devices. The ability to construct, analyze and match full downstream signatures in multiple portions of each individual premise within a node, in addition to the traditional DOCSIS frequencies, means that an operator can develop an accurate map of all the impairments in the upstream and downstream at virtually every component level within that node. Individual premise issues can be isolated and identified, and with the integration of system design prints, network level impairments can be correlated to active or passive components with a fairly high degree of certainty. This allows for impact scaling, and prioritization of impairment resolution with greater precision, as well as improved task management. Since causality can be attributed to the component level with greater accuracy, dispatch of the proper fix agent is more effective and efficient.

4.1. PNM Software Poller

As the DOCSIS specifications have evolved, so has our ability to poll CPE for meaningful data to analyze the performance of the network. Multiple pollers and polling frequencies can be aggregated to form a comprehensive picture of the performance of the network and premises. Using multiple pollers, or polling timing sequences can be useful in managing the demand on the systems by spreading the requests across time, while optimizing which data points are required at specific intervals. For instance, one PNM software programs could poll devices 3 times per day to report raw measurements describing connectivity between the modem and the CMTS, while others collect different data sets like Full Band Capture (FBC) energy in the FM band at more frequent intervals. We describe some of that data collected by the various pollers in the tables below.

Table 1 – Measures Collected by the PNM Poller

Telemetry	Description
Downstream Receive Power	Power (decibels relative to one millivolt, dBmV) received by the cable modem on the downstream DOCSIS channel. For a DOCSIS cable modem to work within the specification, the downstream power level needs to be in the -15 dBmV to + 15 dBmV range.
Upstream Receive Power	Power (in dBmV) received by the CMTS from a specific modem.
Downstream SNR	Signal to Noise Ratio (SNR) refers to the strength of the signal being received relative to the noise on the line. For the cable modem to

	work within the specification, the SNR needs to be at least 23.5dB; a desirable ratio is 30 dB or higher.
Upstream SNR Individual Modem	Similar to Downstream SNR (see above) but for signal sent upstream
Upstream SNR Channel Average	Similar to Downstream SNR (see above) but for the upstream channel average
Upstream Transmit Power	The strength of signal transmitted by the cable modem. Target levels could be between 42 and 50dB.
T3 Timeouts	Number of instances of T3 Timeout (i.e. ranging request retries exhausted). The cable modem has sent 16 Ranging Request (RNG-REQ) messages without receiving a Ranging Response (RNG-RSP) message in reply from the CMTS within 200ms. The cable modem is therefore resetting its cable interface and restarting the registration process. A T3 timeout is typically associated with an upstream impairment.
T4 Timeouts	Number of instances of T4 Timeout (i.e. when the cable modem did not receive a station maintenance opportunity to transmit a Ranging Request (RNG-REQ) message within the T4 timeout period (30 to 35 seconds) after receiving the previous RNG-RSP). The CMTS must provide each CM a periodic ranging opportunity at least once every T4 (30 to 35) seconds. A T4 timeout usually initiates a reboot of the CPE due to lack of communication: the cable modem resets its cable interface and restarts the registration process, reinitializes its mac after T4 seconds have elapsed without receiving a periodic ranging opportunity. This typically indicates an occasional, temporary loss of service, however, if the problem persists, it may indicate possible service outages or maintenance activity on the CMTS. A T4 timeout is typically a downstream problem; if many modems are affected by T4 errors it may indicate a section of the HFC network being affected. T4 timeouts can also be associated with a CMTS that has extremely high usage (e.g. > 95 percent capacity).
Lost Syncs	Number of times the CM lost synchronization with the downstream channel. Discontinuities in the value of this counter can occur at reinitialization of the managed system.
Resets	Total number of resets happened on the CM
System Uptime	Time in seconds (s) that the device has been 'up' and running.
FEC (Forward Error Correction)	Data does not always transmit through cables perfectly, there are often some errors in the bits (1s and 0s); a 0 might be flipped to a 1 or vice versa. To correct those errors, some extra data must be attached to every codeword that goes out (a codeword is just a fixed chunk of data). If it can fix the error, then all is well. However,

	sometimes codewords have too many incorrect bits and the error correction algorithm cannot fix it. If more than 1 percent of the codewords are uncorrectable, we may start to encounter some issues. FECBlks—The total number of FEC blocks (both good and bad) received by all the upstream ports associated with a given downstream. FEC is generally measured using codewords. A codeword is 16 bits with 2 bits for error correction.
Upstream FEC Unerrored	Total number of code word blocks received without any errors. In ideal scenarios it should be equal to total FECBlks
Upstream FEC Corrected	The total number of FEC blocks received by all the upstream ports associated with a given downstream that were slightly corrupted by noise or ingress and that could be corrected and recovered by the FEC algorithm. Number of code words detected as errored, and which have been corrected.
Upstream FEC Uncorrectable	The total number of FEC blocks received by all the upstream ports associated with a given downstream that were so corrupted by noise or ingress that they could not be corrected or recovered by the FEC algorithm. High values imply presence of data loss; this can happen if there is lot of noise in the signal.
Downstream FEC Unerrored	Similar to Upstream FEC (see above)
Downstream FEC Corrected	Similar to Upstream FEC (see above)
Downstream FEC Uncorrectable	Similar to Upstream FEC (see above)
Upstream | Downstream Interface ID	Generally, a device will be connected to the same interface ID over a period of N days; a change to another interface can be an indicator of network-related issues
Upstream ICFR	In-Channel Frequency Response: the max-min variation peak to valley of the response represented in dB. Minimal to no variations are ideal; large variations indicate network integrity issues.
Upstream Ripples	The number of amplitude variations within a 0.75dB threshold.
Upstream TTE	Total Tap Energy upstream from CM expressed in dB
Upstream PMT	Post Main Tap – the Adaptive Equalization (EQ) determined tap that is expending the most energy beyond the main tap. (In the absence of RF impairments, all signals will traverse through the main tap which will have the highest value.)
Upstream PMTE	Post Main Tap Energy – The PMT amplitude expressed in dB. Energy observed in these PMTs is indicative of RF impairments such as micro-reflections being present and the CMTS is activating the equalization gain states in its attempt to compensate for impairments. Energy level is positively correlated to the degree of impairment.
Upstream ETAP1, ETAP2	The EQ tap information for echoes or impedance mismatches beyond the PMT greater than the established threshold. ETAPs provide approximate 'distance' to issue (fault/echo). "Distance to issue" can be inferred as the ETAP number post "main"

	tap (outside premise that serves a cluster of residences) with highest 'energy' multiplied by approximately 85 ft per ETAP increment) [2]
CM Phase	Cable modem signal's measured phase angle deviation before adaptive equalization; frequency carrier wave's phase information is a key component in Quadrature Amplitude Modulation (QAM) technique for transmitting digital data as an analog signal.

4.2. PNM Account and Device Network Analysis Degradations Tool

Comcast has a PNM tool that seeks out account and network degradations. The tool consumes data from multiple pollers and runs the impairment algorithms to examine devices six times per day, to report account and device level degradation issues related to disturbances in the RF spectrum. Using comparative analyses, the tool attempts to isolate impairments related to network, drop, in-home wiring, loose connections, and whole home concerns. The table below lists selected metrics used in the impairment detection.

Table 2 – Measures Analyzed by the Algorithmic Impairment Analysis Tool

Metric	Definition
In-Channel Frequency Response (ICFR)	ICFR is the peak to valley measurement of the amplitude flatness. Ideal ICFR is flat and high ICFR can be caused by impedance mismatches due to loose connectors or other cable integrity issues. The algorithmic tool checks for matches between the ICFR signatures and PMT (Post Main Tap) values to identify a common impairment across neighbors.
Carrier to Interference + Noise Ratio (CINR)	CINR is defined is the range between highest and lowest Modulation Error Ratio (MER). CINR help us capture ingress on non-FBC-capable devices.
Full Band Capture (FBC) FM Noise	Full Band Capture -capable devices can identify if there are FM signals present at the RF input. Loose connectors and other cable integrity issues cause FM signal ingress.

Table 3 – Impairment Locations Identified by the Algorithmic Impairment Analysis Tool

Impairment Location	Description
Solo	The tool cannot isolate the fault to specific location in the home or drop but have checked neighbors and excluded a potential network issue. Fault can be at the device, ground block, drop but is not isolated to that single outlet or device.
Outlet	The tool isolated the fault to the specific outlet that feeds a single device.
Home	The tool isolates the fault to a common point in the home that feeds multiple devices but does not impact all devices in the home.

Drop	The fault is isolated to a common point that feeds all devices in the home and is not a network issue.
Network	The tool has correlated at least one other near neighbor with the same impairment on the same channel(s) indicating an Outside Plant impairment.

4.3. PNM Node-Level Network Degradation Impairment Analysis Tool

Alongside the account- and device-level impairments, there is a tool that consumes both account-level data as well as node-level data. This algorithmic tool that scans the network for RF impairments and reports them continuously and creates event logs indicating a problem and a list of accounts impacted. Events move from soaking to confirmed stages and vary by severity. Events include outages, service call alerts, plant faults, spectrum power analysis events, and upstream CMTS port analysis events.

Table 4 – Node-Level Impairments Identified by the Algorithmic Analysis Tool

Event Type	Description
Outage	Event generated when 4 or more devices that share a common problem are offline. Criteria area adjusted for multi-dwelling units.
Power outage	Event is declared if 4 or more devices appear to be a part of a commercial power outage, identified by devices on battery backup.
Service call event alert	Event is generated if the percentage of customers with a scheduled trouble call on the node exceeds a threshold
Plant Fault	Events are generated when a proportion of customers within the node violate the thresholds for the RF Levels. Problems include out of range upstream transmit power, upstream or downstream SNR/MER, downstream receive power.
Spectrum Power Analysis Events	Events are generated by looking for a cluster account matching a specific RF spectrum signature. Signatures include suckouts (a frequency-specific collapse in the RF energy caused by a short distance and high energy reflection), waves (a sinusoidal pattern detecting multiple peaks to nulls in the spectrum 4.5 dB or greater), and tilts (a delta threshold between upstream transmit level in the highest and lowest frequency channels that exceeds 3dB)
Upstream CMTS Port Events	Events are generated by monitoring utilization of upstream channels, failed and correctable FEC errored rate, number of upstream channels in use, and channel width and modulation.

We discuss how the inputs from the remote network telemetry are transformed into the model features in section 6.1.

5. Refer to Maintenace (RTM) Field Operations

When customers encounter service issues, they can reach out for help through various channels, like IVR (Interactive Voice Response), call or chat agents and the Xfinity Assistant interactive chat application. If a chatbot application detects issues that cannot be solved by system refresh or restart, it will route subscribers to a chat or a call agent. Call and chat agents use Interactive Troubleshooting Guides (ITGs) to localize the problem and offer the most effective solution. If an agent identifies an impairment in the network, they will schedule a service technician to visit customer's home and continue troubleshooting on site.

Ventriglia et al [3] describe a workflow for on-site troubleshooting. On site, technicians use field tools to take measurements at different points in the drop network. A technician might start with taking measurements at the ground block closest to the customers' premise. If these measurements are within acceptable specs, the technician will move to investigate impairments inside the customers' home. If a fault is found in the grounding block, the technician will continue to investigate issues at the tap. If no issues are found at the tap, the technician will do the work to replace the drop and perform a premise health test to confirm that customers' issues are resolved. If an impairment is found at the tap, the technician will escalate the problem to the line tech team by creating an RTM request. Issues are triaged by field operations teams and if they are confirmed as a plant impairment, a line technician job is scheduled. Field operations teams monitor non-severe issues, identify neighborhood impact, and send a line technician once a neighborhood impact is confirmed. When the standard technician finds a network issue, a cost is incurred to send both a standard technician and a line technician. Moreover, we estimate that a subscriber might wait up to two days for a standard technician appointment and then possibly another day for the line technician appointment. This historically has been the repair model when a single customer reports a service issue that is ultimately attributed to an impairment in the outside plant network.

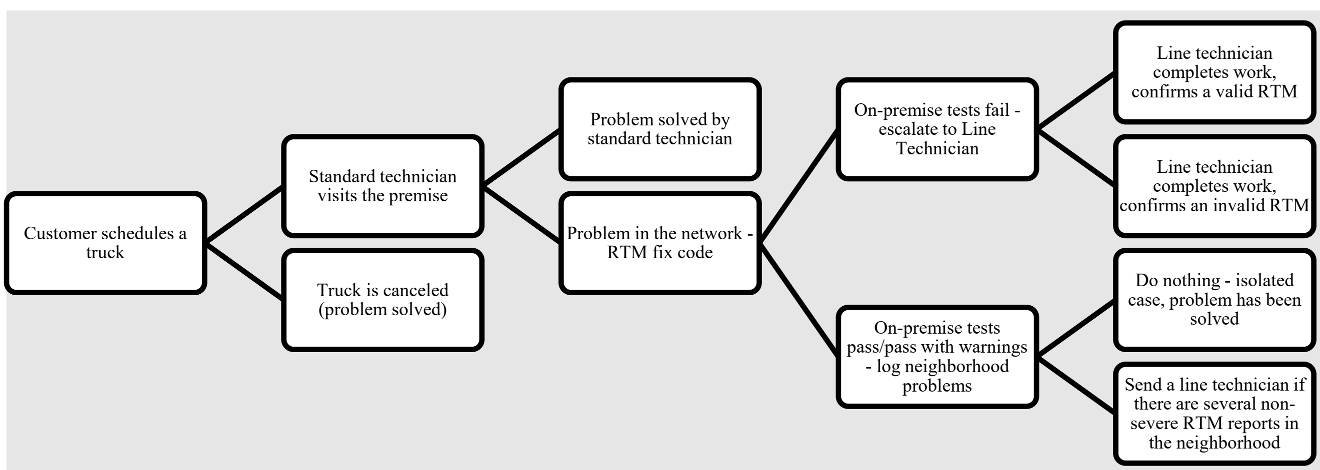

Figure 2 – Refer to Maintenance Field Operations

6. Refer to Maintenance Machine Learning Model

We propose a workflow where a predictive model will recognize that the issue will be an RTM problem, with high probability. In this case, we will not send a standard technician, but schedule a line technician instead. If the model predicts a low RTM probability, we might message the standard technician that the issue is an unlikely RTM.

We have developed a machine learning classification model to identify standard technician visits that will result in the RTM escalation that is also confirmed as “valid”. The model is trained on the set of data comprised of standard technician visits. The model learns the target labels derived from the line technicians the “valid RTM” label, whereas all other codes are included as negative labels. We exclude trouble calls where technician could not access the premise, or no one was home.

The model has access to two sets of features: telemetry collected by pollers for all subscribers in the company footprint and information collected ad-hoc for subscribers who reported issues during the call with an agent.

We used a two-step modeling approach, which we refer to as *tier 1* and *tier 2*, to make the final prediction. Our *tier 1* model is trained on all the available telemetry information to determine the probability of RTM among the subscribers’ neighbors. This is then used to evaluate the probability of RTM for the subscriber during the day prior to the day of call. Subsequently, our *tier 2* model uses features collected ad hoc during the call with an agent, archived telemetry features, including the *tier 1* derived probability of RTM during the previous days as well as the probability of RTM among neighbors.

Following sections offer details of model development within each stage.

6.1. Feature Engineering

6.1.1. DOCSIS Telemetry Features

The model uses features from PNM telemetry sources that collect and aggregate DOCSIS measurements. These sources poll and analyze the state of the network, looking for outages and impairments. Other systems collect network data and do the fault segmentation analysis.

We batch-process features from the PNM sources once per day. Each source has a varying aggregation window. Data from the pollers reporting connectivity between modem and CMTS and account degradation algorithm tools arrive 3 to 6 times per device per day.

For the tool reporting account-level degradation, we used a one-day lookback window and aggregated the data across devices and polls reporting the proportion of polls with a certain impairment. Aggregated features include a potential filter indicator, number of correlated accounts found by the algorithm, FBC FM Noise and ICFR impairments.

The PNM tool reporting on connectivity between the model and CMTS offers measurement and count data. Devices report cumulative counts of T3 and T4 timeouts, lost syncs, resets, FEC corrected, uncorrectable, and total codewords until they are rebooted. We transform count features to get the incremental increase since the last poll unless there has been a reset, otherwise, we take the feature value as is. Measurement features include receive and transmit power, SNR, TTE, PMT, PMTE, CM phase, upstream ripples, and ETAP1 and ETAP2 measures. We further transform FEC errored and corrected codeword counts as percentage of total errored and corrected codewords. Then, we proceed with data aggregations: we take averages for features across all polls for all devices attached to the account during the day prior to the model evaluation. Finally, we examine data across four days prior to the model evaluation to estimate standard deviations for the features.

We collect data from the node network impairment algorithm for all events created one day prior to the model evaluation tied to the subscriber who requested the truck. We estimate the number of accounts impacted by the same events and identify events by type: plant issues (ICPF plant faults, upstream and downstream FEC), wave, and suckout events. We count the number of events reported for each account.

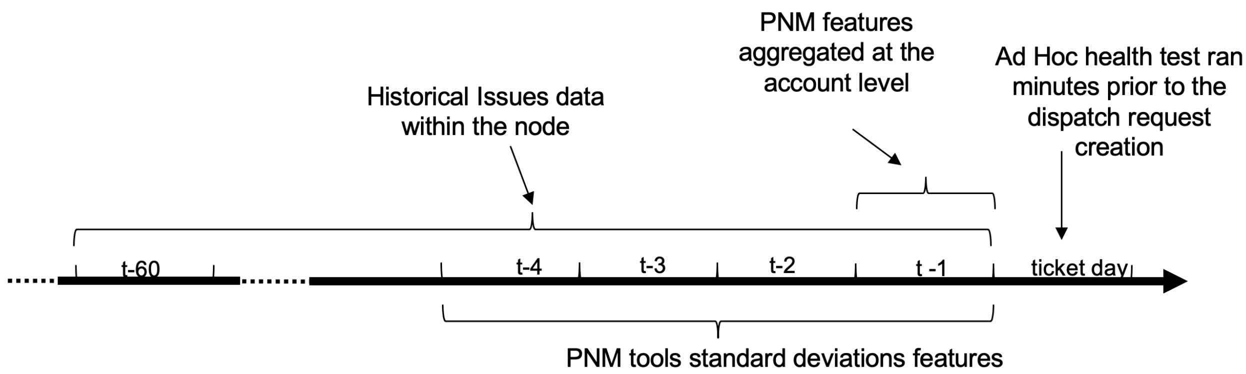

Figure 3 – Feature Extraction Timeline

6.1.2. *Historical Issues Reported by Technicians*

Another data source added to the model are historical rates of different types of issues reported by the technicians in the same nodes in the previous 60 days. Issues include construction and underground crew escalations, drop, ground block, and prior RTM escalations. We include these features as rolling aggregations for problems 60 days prior to the model evaluation.

6.1.3. *Ad Hoc Remote Telemetry Health Checks*

Agents use the remote Application Programming Interface (API) to run ad hoc telemetry checks using the capabilities of polling tools to collect the most recent data about the customers' network. Data from PNM pollers are collected 3 to 6 times per day and lag in time due to processing and data availability delays. Ad hoc telemetry health tests offer data from the same pollers but are retrieved on demand. In addition to the raw data, the tests also report pass/fail/pass-with-warnings analysis and offer heuristics for interpreting the raw data ranges (such as labels for failures in each DOCSIS analysis section). We transform raw data with similar aggregations described for PNM tools earlier and add labels and heuristics data as one-hot-encoded or ordinal (based on severity levels) features.

6.1.4. *Topology Aggregations*

Node topology was instrumental in creating features that detect network impairments. Node topology maps describe how equipment such as amps, taps, splitters, and cables are connected from the node to customers' homes. Data are collected by calling on the geographic information system API with a backend database to execute spatial intersect queries to identify equipment related to the locations supported by the Comcast footprint. Data are stored as a directional graph representing paths from the node to the tap, along with coordinate and equipment details.

Node topology maps were used to aggregate and average measurements, to provide a wider view of the network events surrounding a customer when the service call was being scheduled. One type of aggregation is at the parent level, which averages the telemetry over all customers who are immediate neighbors in the graph. Another aggregation is a comparison of parent measurements to those of parent's neighbors. The idea is that if a single piece of network equipment is broken, we can observe the impact of this by comparing all the customers who depend on it (and thus have impaired service) with customers who do not depend on it but are otherwise similar because they share the same upstream network components. We also identified locations on the graph where amplifiers are located and estimated the average and standard deviation of measures across amplifiers on the path from customers' home to the node.

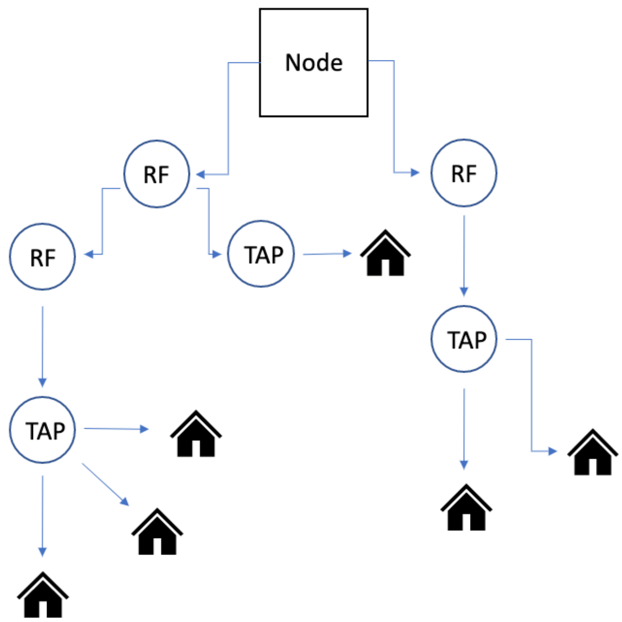

Figure 4 – Node Topology Diagram

Figure 5 provides an overview of the data, feature engineering and integration in our two-step modeling approach.

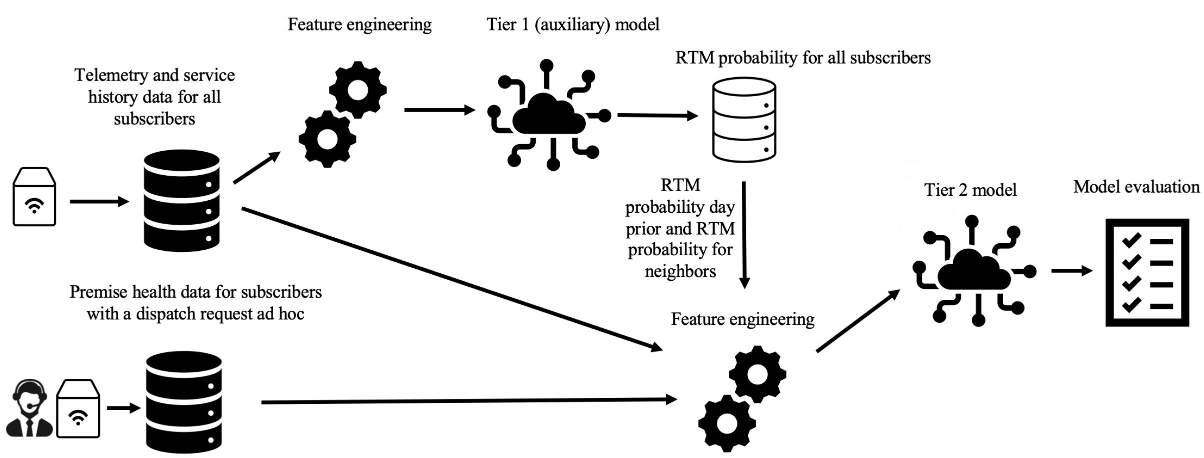

Figure 5 – Model Training Process

6.2. Model Evaluation

6.2.1. Model Performance

Our best-performing models is trained with an open-source XGBoost classifier (Tianqi and Carlos, 2016 [4]). The classifier has been calibrated for best hyper-parameters and uses 500 estimators to build trees with maximum depth 2. 239 features are used as model inputs in *tier 1* model and 375 features are used as model inputs in *tier 2* model. Our target label has a very small prevalence, resulting in the class imbalance problem. To assist the models to learn from a data set with class imbalance, we used the class-weighted version of the model to help remediate the class imbalance.

During the training stage, we separated data into training and testing sets, using the time of ticket creation to separate testing set of observations. Further, *tier 1* and *tier 2* models are trained on separate sets of data. In the latest iteration, tier 1 model was trained on 2,231,591 observations collected during December 2020 through March 2021.Tier 2 model was trained on 1,021,398 observations collected during April through June. During evaluation, tier 1 model is tested on full available out-of-sample data (April through August 2021), however, here we quote performance numbers for both models tested on the same data with 587,975 observations, collected during July through August 2021.

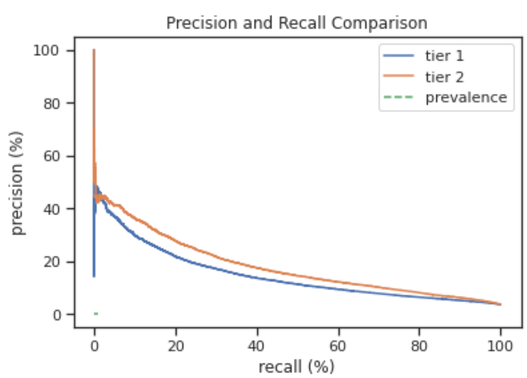

Figure 6 – Precision Recall Curve on Testing Data

The *tier 1* model is trained with telemetry data available for all subscribers. Calibrated at 5 percent recall, the model achieves 37.2 percent precision. Model performs 10 times better compared to the random guess.

The *tier 2* model is trained with combined telemetry, RTM probability of neighbors and subscriber prior day RTM probability, and ad-hoc health test features. Calibrated at 5 percent recall, the model achieves 41 percent precision: this means that the model calibrated to return at least 5 percent of true RTM trucks, can accurately identify 41 percent of cases. Model performs 11.2 times better compared to the average prevalence (random guess).

While Model stacking increases complexity of the model training process, it helps us to increase model precision by 3.5%. Figure below shows the precision and recall curve for the *tier* 1 and tier 2 models evaluated on the test data.

6.2.2. Exploring Feature Importance

We used normalized model weight to explore top features contributing to the model importance. Following graphs show features ranked and identified by source.

In *tier 1* model, features extracted from the account degradation analysis tool make up the majority of top features. Many of the features contributing to the model are based on the topology estimates – you can identify them by "parent" (tap aggregations) or "AMP" (amplifier aggregations) flags.

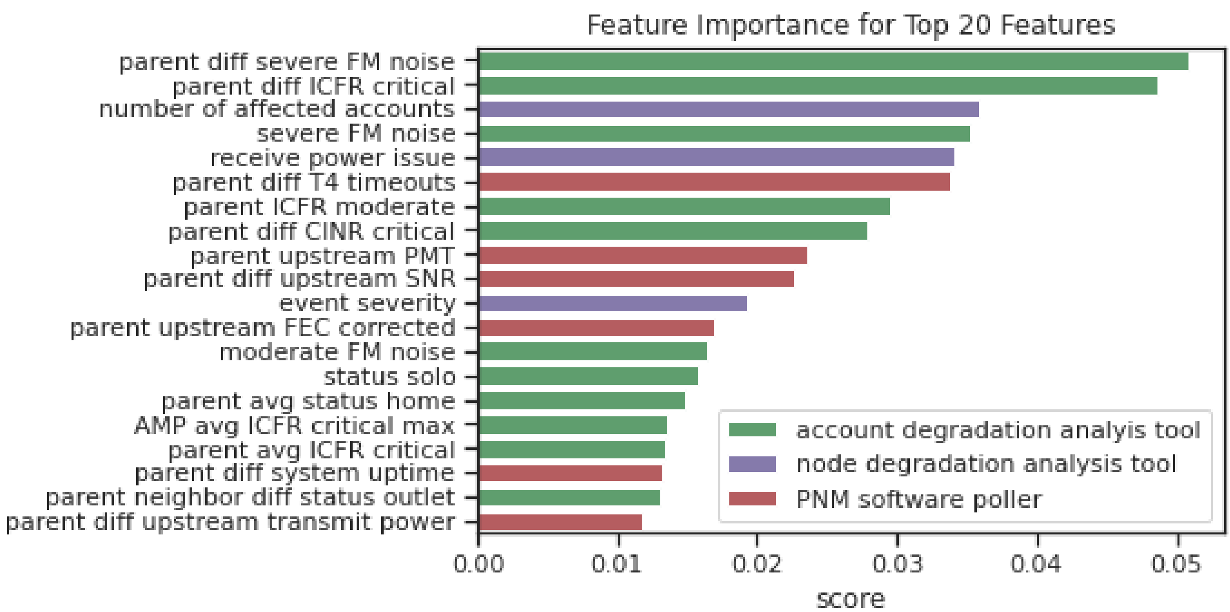

Figure 7 – Feature Importance for Tier 1 Model

In *tier 2* model, features extracted from the premise health check tool make up the majority of top features – it is expected as these features are extracted during the call with the agent and represent the most recent state of the device performance. The top 2 features contributing to the model come from the auxiliary tier one model and indicate probability of RTM 3 days prior to the call. The importance of these features leads us to believe that network problems might impact customers in a matter of days prior to the call. We also explore the directionality of features for *tier 2* model.

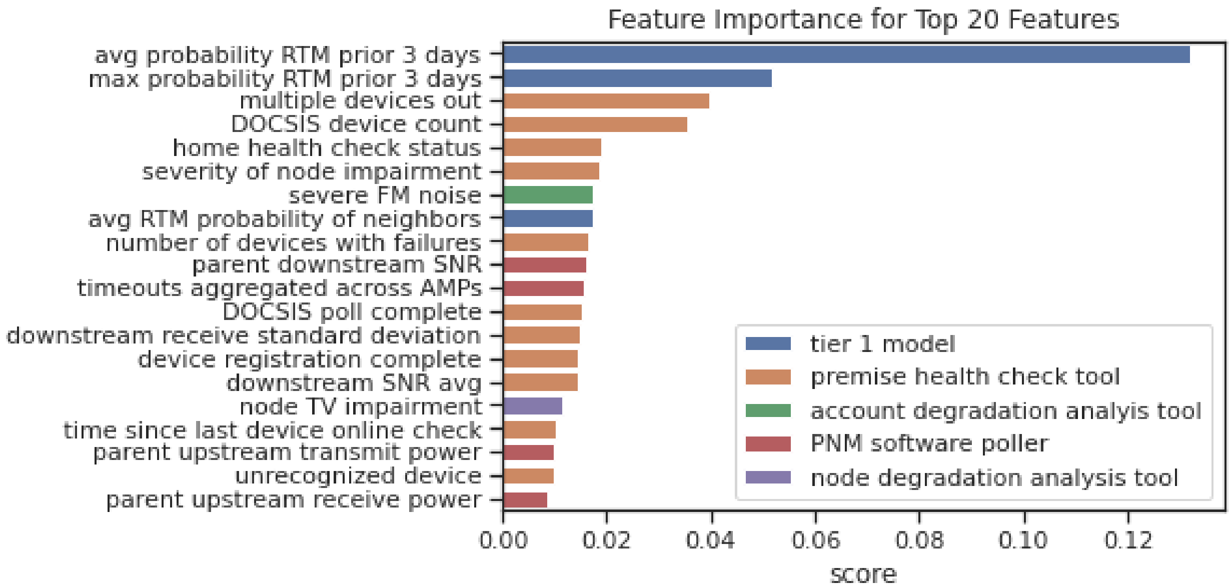

Figure 8 – Feature Importance for Tier 2 Model

Features positively correlated with the RTM outcome are number of device failures identified by health check, probability of RTM days prior, high downstream receive standard deviation, failed (high value of) home health check status, multiple devices out/offline, number of events with a TV-related impairment evident in the node, severity of node impairment, severe FM noise event incidence, and timeouts aggregated across amplifiers. For these features, higher values indicate a higher chance of RTM.

Features negatively correlated with RTM probability are time since last device online check/registration, parent upstream receive power, complete DOCSIS polls, found unrecognized devices, DOCSIS devices on account, complete registrations, downstream SNR, and parent downstream SNR. For these features, lower values indicate a higher chance of RTM.

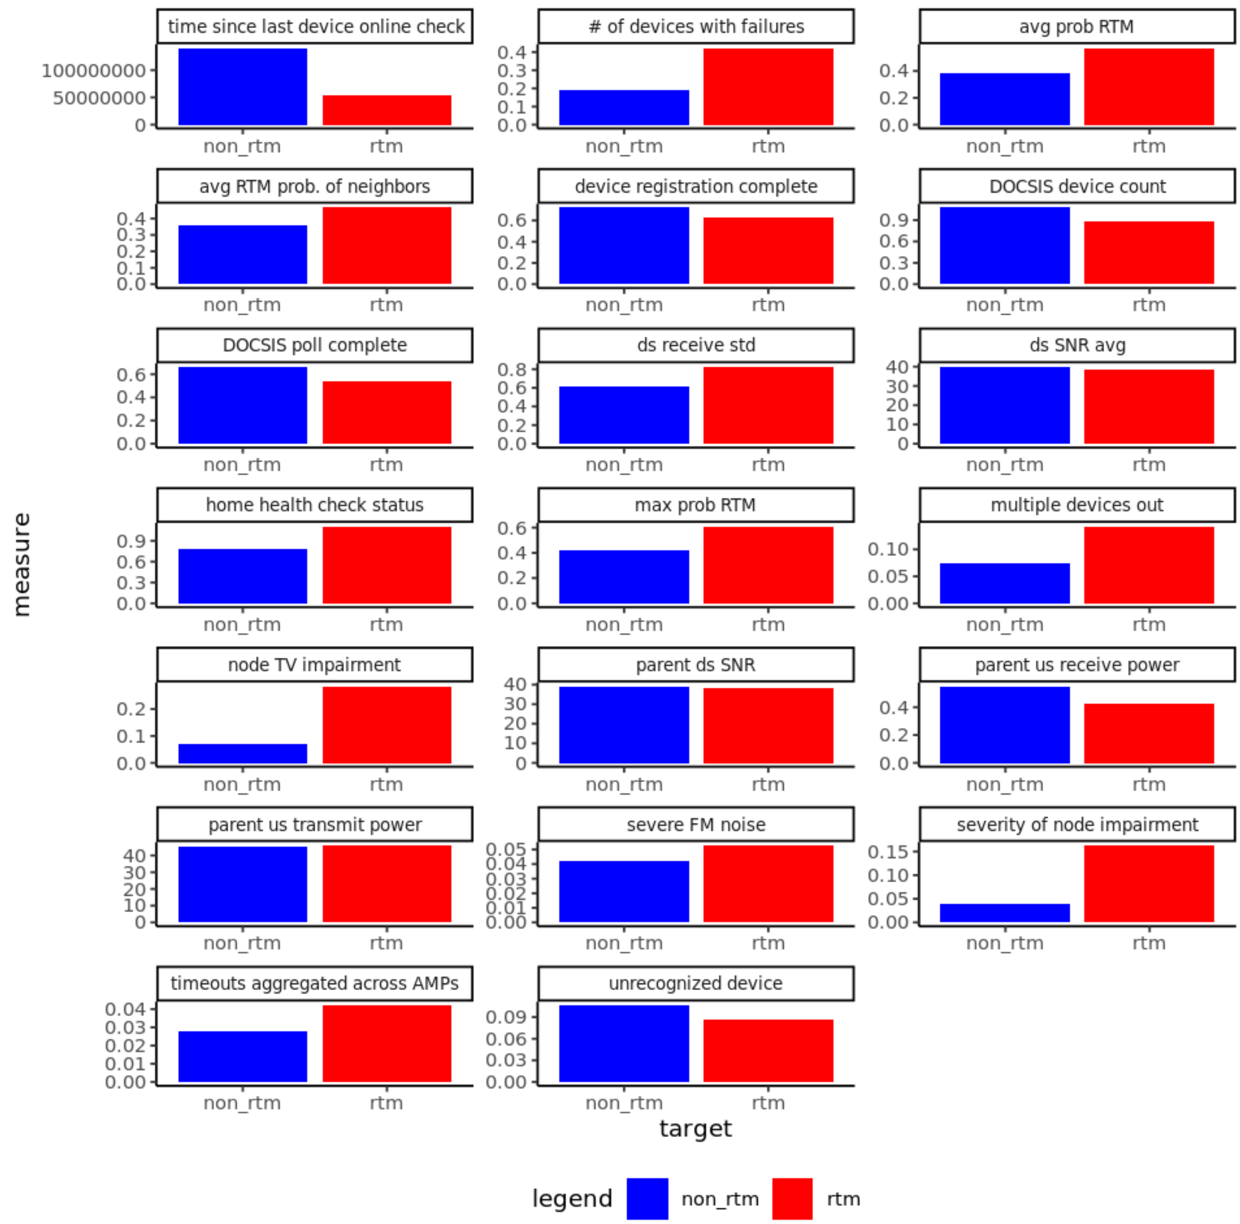

Figure 9 – Feature Exploration for Tier 2 Model

6.3. Cost Calculation

To assess the financial impact of the RTM model, we constructed a cost analysis comparing the business as usual (BAU) process to our model (RTM). The goal of this analysis is to identify the minimum precision required by the model to incur a net zero savings relative to BAU. This allows us to know when our model precision exceeds the break-even point before deploying in a live environment.

We assess the net savings relative to BAU when the RTM model makes a positive prediction (i.e. when the model recommends sending a line technician instead of a fulfillment truck). Note that the net savings for a negative prediction is effectively zero since a negative prediction defaults to the BAU protocol.

The cost analysis presented here also assumes that the model will make recommendations without any subject matter expert (SME) interventions. Asking SMEs to review model predictions, thereby utilizing the "human-in-the-loop" framework, improves the model precision but its impact to the overall cost is yet to be estimated.

Running a machine learning model incurs a fixed cost of computing, data storage, and engineering resources. Additional costs might include if we use a human-in-the loop model and each prediction is reviewed by an expert. The cost to run a model would be a fixed cost per prediction.

There are two approaches we used to assess the net savings incurred by our model. First, we describe in more detail the simpler of the two approaches. *Customer-level approach* estimates the net savings by taking into consideration only the customer with the model prediction in question. The latter approach, referred to as the *neighborhood approach*, additionally takes into consideration the impact of sending a line technician to the customer's neighborhood. Sending out technicians early and preventing more fulfillment trucks from going out into the neighborhood can generate additional savings. Let's first denote the following costs:

Table 5 – Cost Model Definitions

Cost of fulfillment truck	F
Cost of line truck	L, where L > F
Operations and research cost per prediction	R

6.3.1. Customer-Level Cost Calculation

In a true positive scenario where the RTM model correctly predicts that a line technician is required, the RTM model incurs a cost of L and the BAU model incurs a cost of sending both trucks, $L + F$. Therefore, the net savings for a true positive is

$$(L + F) - L - R = F - R$$

In a false positive, the RTM model incorrectly recommends a line technician and we will assume that in this case, we would need to send a fulfillment truck after incorrectly sending a line technician, costing the company $L + F$. We note here that this is a very conservative estimate as we can expect that sometimes a line technician can resolve issues that normally require a fulfillment truck (e.g. when an issue exists outside the home). For a false positive, the BAU's cost would simply be F. Let's also assume that sometimes customers cancel their scheduled trucks, and we will denote this cancellation rate as Y. Then, the net savings for a false positive is

$$F * (1 - Y) - (L + F) * (1 - Y) - R = -L(1 - Y) - R$$

We define precision as the proportion of predictions the model made correctly. Using these two estimates, the net savings, S_{net}, of the RTM model with precision P is given by

$$((F - R) * P) - (L * (1 - Y) - R) * (1 - P) \ = \ S_{net}$$

By setting $S_{net} \ = 0$, we can solve for precision P to get the precision required to break even with BAU's costs, which we denote as P_{even}. In customer-level approach,

$$P_{even} \ = \ \frac{L(1 - Y) - R}{F - 2R + L(1 - Y)}$$

Graph below illustrates the relationship between savings and the model precision. Actual estimates are omitted to comply with our company's data reporting policies. We offer the direction of the positive savings and show the precision required to achieve the break-even point.

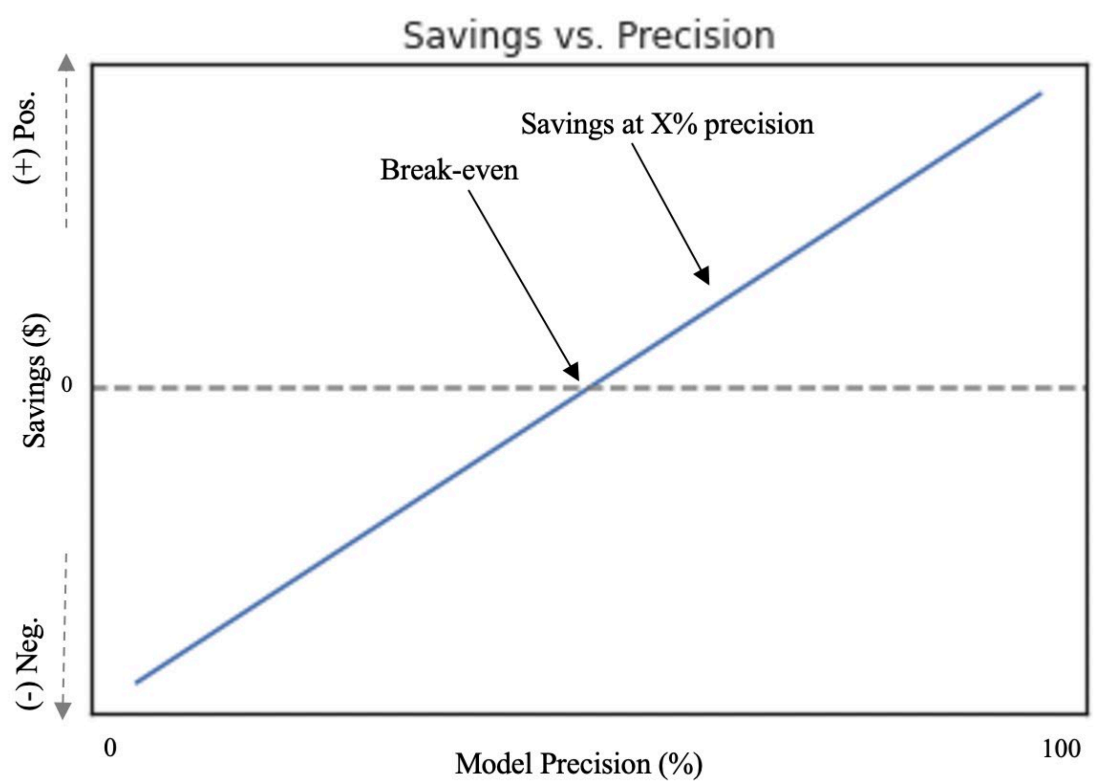

Figure 10 – Precision vs. Savings Curve Illustration

6.3.2. Neighborhood-Level Cost Calculation

In the neighborhood-level analysis we also consider the RTM model's impact to the neighborhood. By sending out line technicians early, there is a positive impact to other customers in the neighborhood suffering from internet issues that are predicted to also be RTM issues. In this scenario, if there exists a customer C_A with an RTM prediction, then we will delay sending out a fulfillment truck to C_A's neighbor, C_B, if the model also recommends a line technician for customer C_B within 2 days of C_A's prediction. In this analysis, we define two people as neighbors if they not only share the same CMTS node but also all the devices between the CMTS node and their respective homes.

In this case, we consider the true positive and false positive costs from customer-level calculation for the original customer C_A and we *additionally* consider the cost of a true positive and false positive to a neighbor C_B. A true positive for neighbor C_B occurs when the model correctly predicts an RTM for the original customer C_A *and* correctly predicts an RTM for customer C_B. In this case the true positive net savings for customer C_B is the cost of sending out a fulfillment truck needlessly, F.

A false positive for customer C_B occurs when we delay the customer from receiving a fulfillment truck when we incorrectly predict an RTM for C_A *or* C_B. The cost here is much less severe than only customer-level approach, as it is mostly a delayed response and possibly a second call to an agent made by the neighbor. If we denote the cost of contacting an agent as A and we assume that Z percent of the time, customer C_B will opt to call an agent instead of using an automated scheduling service, then the cost the model incurs for being wrong with customer C_B is $A * Z$. Because both the neighborhood sizes vary and the number of trouble calls made in a neighborhood vary across customers/regions, we used the test data to generate the net savings of our model as a function of the model's precision. We found that this neighborhood analysis decreases the required precision to break even, suggesting that implementing this delay service will have an additional positive financial impact.

6.4. Model Deployment

6.4.1. Model Deployment for Proof of Concept (POC)

We explored a "soft" deployment of the model for selected markets during our POC stage of development. Our goal was to obtain buy-in from our colleagues and test the model validity without incurring negative costs to the business or impacts to the customer experience.

During the POC we worked with the "human-in-the-loop" strategy. In this process, our program evaluated dispatch requests created by agents every hour and flagged "highly predictive" RTM candidates identified by the model. We then sent emails with the predictions to the regional subject matter experts (SMEs) for triage using AWS SES (Simple Email Service) email automation and posted in the Microsoft Teams channel. Emails were sent right after we received dispatch request logs to simulate a near real-time processing and prevent SMEs access to the truck outcomes (so that their judgement won't be "poisoned" by the truck outcome knowledge).

One of the participating markets explored dispatching a line technician the day a prediction was deemed valid by the SME, with the intent of remediating any customer pain from plant impairments prior to the standard technician visit. The majority of the market SME's evaluated network telemetry via desk top tools, logged notes and then validated the label (whether RTM is needed or not). Of note: the markets opted not to cancel the standard technician visit and schedule a line technician in their place until the predictions provided a higher degree of confidence. We have collected over 4,500 observations from the POC stage. SMEs evaluation of the model predictions resulted in 48 percent model precision. One of reasons for higher precision during the trial could be that evaluations were made on the dispatch request and not on the confirmed service calls. On average, 35 percent of scheduled service calls are canceled.

A debrief with the involved SMEs helped us to discover additional data sources, improve our feature engineering process, and identify false positive cases.

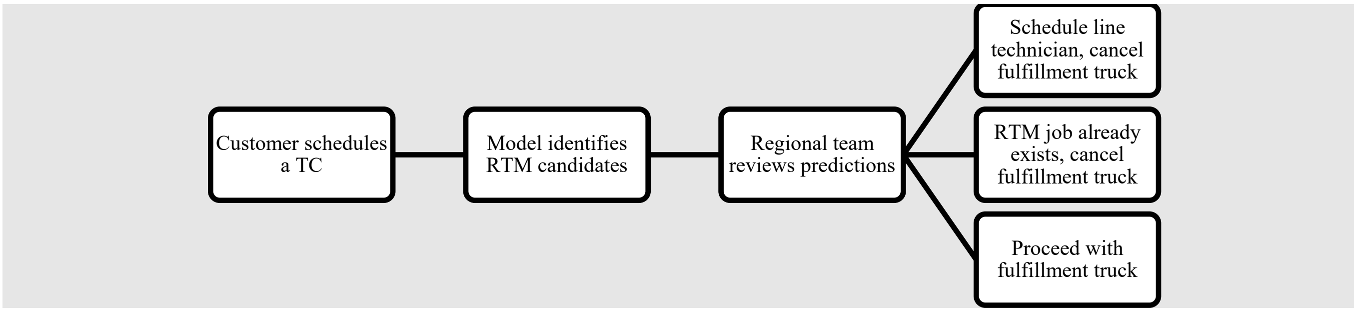

Figure 11 – POC Stage Model Deployment Workflow

6.4.2. Production Model Deployment

Figure 10, below, illustrates how we plan to deploy the RTM model into production. We will rely on the Redis (Remote Dictionary Server) [5], a fast, in-memory, open-source key-value data store to host the data. Our internal platform will host the model API. Model will consume two sets of features described earlier: payload from the ad-hoc premise health check passed by the agent making the API call, and features derived from the archived telemetry and *tier 1* model input stored in Redis feature store. Model engine will process features and output predictions to pass to the call agent. Subsequently features used and model result will be logged in AWS S3 (Amazon Web Services Simple Storage Service).

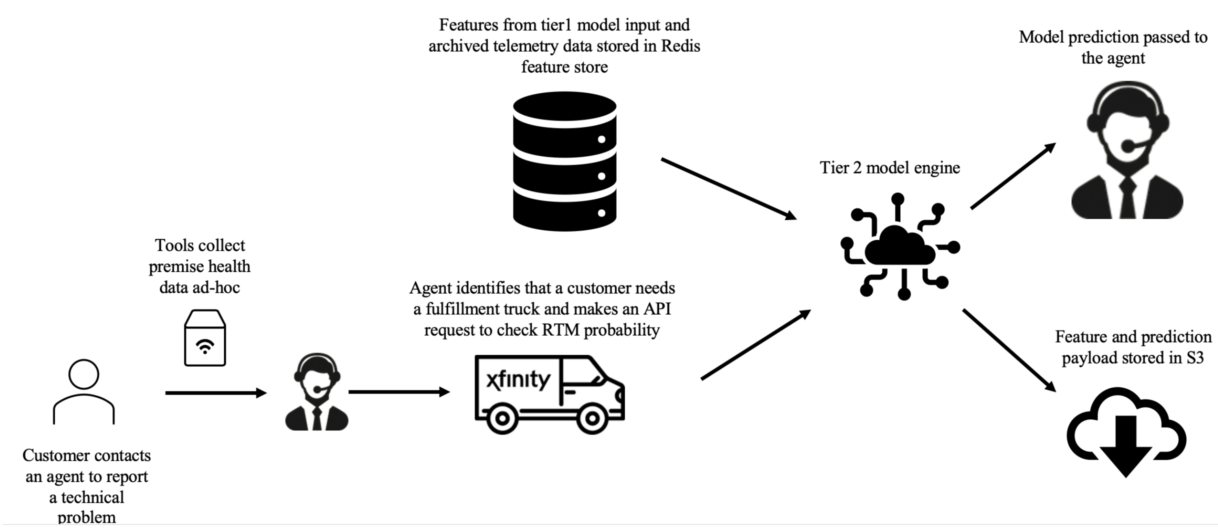

Figure 12 – Model Deployment Diagram

7. History of Model Development and Future Work

Network telemetry and topology features are used across multiple use-cases across the organizations. Originally, a set of features derived from the network telemetry was developed and applied to the classification problem predicting whether technician will only need to do work outside of the home. This model was implemented during COVID-19 business operations and helped us reduce the interactions between customers and technicians. We discuss this model in the Vengriglia et al [2] paper mentioned earlier.

The RTM model has been developed based on the telemetry features and benefited from RTM-specific features derived from topology and RTM probability of neighbors. During the development of the model, we have received feedback from the field operations colleagues and redefined the target label for the model. Originally, we trained our model to recognize the RTM fix code entered by the standard technicians. We then learned that regional teams put the requests through a triage process and only send a line technician for severe service disruptions only, taking other steps to help customers with non-severe service disruptions. Receiving this feedback, we have redefined our target label and updated the model to learn that the standard technician recommended an RTM, field operations triaged the case to send a line technician, and the line technician confirmed that the referral vas valid.

Our model development process does not stop with this paper or the model deployment, which is an early marker in the overall trajectory of ML and network operations. We have worked and will continue to work with experts to identify several telemetry sources that can help us improve model precision. First, we have learned that upstream interface problems can help us detect issues impacting multiple customers. Second, we have also acquired and plan to integrate packet data, hoping that this source can help us detect intermittent issues better.

8. Conclusions

In our paper, we described the complexity of the HFC network and the intelligent systems that collect measurements to monitor the health of the plant. While the telemetry measurements collected from the plant help us detect the symptoms of impairments, the machine learning approach is a predictive approach to plant maintenance and can enable us to dispatch line technicians to the right area in the network and remove problems more efficiently. The implementation of the Refer to Maintenance (RTM) with Machine Learning model would help us move from a reactive to a predictive maintenance approach. In addition, an ML approach can identify and uncover intermittent issues that are not always prioritized (or able to be prioritized) for line technician work, and can deteriorate the customer experience.

Adopting a machine learning approach will help us to continuously deliver an improved customer experience and streamline efficiencies across the service pipeline. Instead of multiple visits by different technicians, our customers could have their issues resolved faster with a single visit by the right technician. Further, we expect that a corrected network issue will have a positive impact on the whole neighborhood, delivering an improved experience to all customers in the area and thus reducing contact rates and unnecessary truck rolls.

We have described how cost analysis can help determine the appropriate model precision to achieve before deployment. Using this approach, we can deploy the model that is tuned to achieve an outcome beneficial to the business. In summary, the Machine Learning approach is a game changer that will help us to achieve continuous operational transformation. By identifying and prioritizing network issues, we can deliver the right technician at the right time and improve customer experience.

9. Disclaimers

We collect, store, and use all data in accordance with our privacy disclosures to users and applicable laws.

Abbreviations

API	Application Programming Interface
AWS	Amazon Web Services
BAU	Business as Usual
CATV	Community Antenna Television
CM	Customer Modem
CMTS	Cable Modem Termination System
CPE	Customer Premise Equipment
dBmV	decibels relative to one millivolt
DOCSIS	Data Over Cable Service Interface Specification
ETAP	Electrical Transient Analyzer Problem
FBC	Full Band Capture
FEC	Forward Error Correction
HFC	Hybrid Fiber Coax
ICFR	In-Channel Frequency Response
ISP	Inside Plant
ITG	Interactive Troubleshooting Guide
IVR	Interactive Voice Response
MDU	Multi-Dwelling Units
ML	Machine Learning
OSP	Outside Plant
PMT	Post Main Tap
PMTE	Post Main Tap Energy
PNM	Proactive Network Maintenance
POC	Proof of Concept
RF	Radio Frequency
RNN	Recurrent Neural Network
RTM	Refer to Maintenance
S3	Simple Storage Service
SES	Simple Email Service
SME	Subject Matter Expert
SNR	Signal to Noise Ratio

Bibliography & References

[1] SCTE Expo 2016 - A Comprehensive Case Study of Proactive Network Maintenance; Wolcott et al

[2] VIAVI 2015 - Understanding Equalizers, Pre-Equalizers, and Their Use in Localizing and Troubleshooting HFC Issues; Fenton, Mike and Jump, Larry

[3] SCTE Expo 2020 – Training Machines to Learn from Signal Meter Readings; Ventriglia et al

[4] XGBoost: A Scalable Tree Boosting System, March 9, 2016; Chen, Tianqi and Guestrin, Carlos

[5] Redis, Documentation retrieved from https://redis.io/