Training Machines to Learn From Signal Meter Readings

A Case Study from Comcast

A Technical Paper prepared for SCTE•ISBE by

Gary Ventriglia
Sr. Principal Engineer
Comcast Corporation
1800 Arch St, Philadelphia PA 19103
484-846-0467
Gary_Ventriglia@comcast.com

Jack Birnbaum
Vice President, Customer Experience
Comcast Corporation
1800 Arch St, Philadelphia PA 19103
215-286-8057
Jack_Birnbaum@cable.comcast.com

Robert Gonsalves
Sr Director, Product Dev & Engineering
Comcast Corporation
1401 Wynkoop Suite 300, Denver CO 80202
720-512-3642
Robert_Gonsalves@cable.comcast.com

Anastasia Vishnyakova
Engineer 3, Machine Learning
Comcast Corporation
1800 Arch Street, Philadelphia PA 19103
267-260-3277
Anastasia_Vishnyakova@comcast.com

Michael Kreisel
Principal Researcher, Machine Learning
Comcast Corporation
1110 Vermont Ave NW, Washington, DC 20005
202-524-5068
Michael_Kreisel@comcast.com

© 2020 SCTE•ISBE and NCTA. All rights reserved.

Larry Wolcott
Comcast Fellow of Engineering
Comcast Corporation
1401 Wynkoop Suite 300, Denver CO 80202
720-512-3643
Larry_Wolcott@cable.comcast.com

© 2020 SCTE•ISBE and NCTA. All rights reserved.

Table of Contents

List of Figures

© 2020 SCTE•ISBE and NCTA. All rights reserved.

List of Tables

© 2020 SCTE•ISBE and NCTA. All rights reserved.

1. Introduction

When informed by vast amounts of network performance information, identifying radio frequency (RF) problems with the Data-Over-Cable Service Interface Specifications (DOCSIS) isn't that hard. However, determining if the problems are inside or outside the home can be difficult. This is a decades-old problem, with hopes often pinned on the elusive promise of artificial intelligence (AI) or machine learning (ML) to help. The challenge that many data scientists will tell you is that having good training data is critical. The lack of a reliable feedback loop to establish cause-and-effect often results in poorly trained machines.

A significant amount of time and resources has been poured into remote diagnostic tools to identify plant problems. Those tools historically have been segmented, specialized, and tuned to evaluate singular aspects of the RF health – for example, receive modulation error ratio (RxMER) and forward error correction (FEC). Once a problem is identified, determining if its source is in a customer's home, drop or tap has historically been left to technicians, to provide feedback about what they found. The feedback mechanisms typically involve selecting a code or result and updating the work order when it's complete.

With the COVID-19 pandemic starting mid-March 2020, the rapid development of an "outside network check" provided an opportunity to gather better features and labels. With a renewed desire to keep technicians and customers isolated, the team is exploring new ML/AI models. These new models are trained to use cloud-based RF measurements. These measurements include remote telemetry from DOCSIS devices, and other equipment logged by collection systems. Another set of measurements is taken at the tap and ground block, finally offering a way to segment the network and train the machines differently. The authors review the outcome of this fascinating exercise currently under way, as this paper is being written, in the summer of 2020.

2. Acknowledgements

Much of the material within this document was obtained from interviews and written contributions of many Comcast leaders and other industry experts. The authors would like to thank David Monnerat, Gary Schwin, Justin Menard, Shawn Hughes, Patrick McDonald, Miles Pellegrini, Marty Marcinczyk, Jan Neumann, Rama Mahajanam, May Merkle-Tan, and Fan Liu from Comcast. We thank May Merkle-Tan for her contributions to the early stages of the outside-of-home model developments and for setting up and supporting the data pipelines required in model deployments. You all collectively represent a large cross-section of unique perspectives and expertise. We also recognize Brady Volpe of the Volpe Firm and his award-winning products and services. Thank you, Brady, for your ongoing industry leadership and inspiration. Finally, we express gratitude to our distinguished colleagues Seamus Gallagher, Dwain Kelly and Paul Kelly from Liberty Global. Your global contributions and experience are invaluable to our industry.

3. A Brief History of RF Troubleshooting

For well over 40 years, cable field technicians have relied on signal level meters (SLMs) to take measurements that help determine a proper course of repair. Once alerted and dispatched to a problem, our technicians have been trained to use a divide-and-conquer approach for troubleshooting, and ultimately repairing the issue. These meters would usually be used to take measurements of RF signals to-and-from the customer at different locations on the network. They might start inside the customer location, to validate the service at the location of the customer's equipment, then work their way towards the network to determine where the problem begins or ends. There are other processes where technicians

© 2020 SCTE•ISBE and NCTA. All rights reserved.

might start at the network tap and work their way towards the customer equipment. Of course, these are basic processes which traditionally made it very difficult to verify if, or how they were followed.

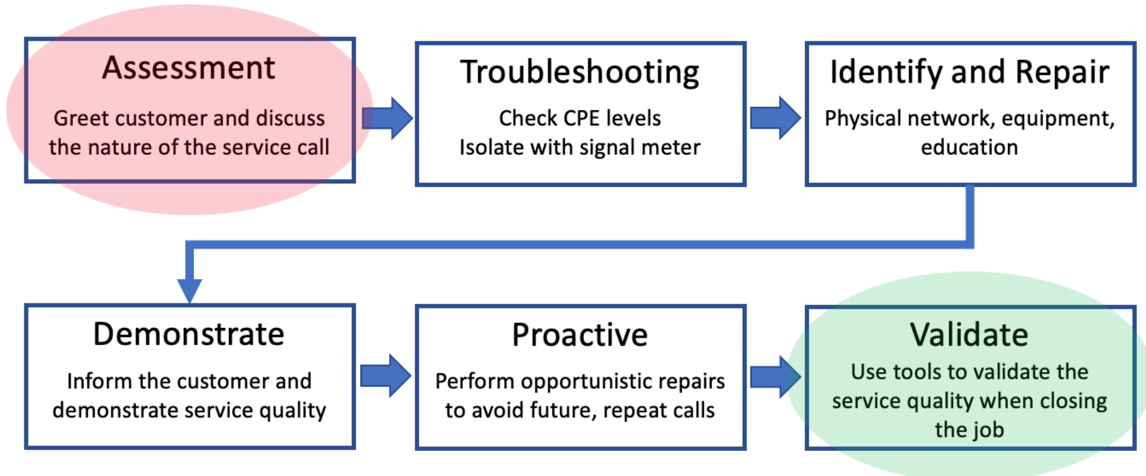

Figure 1 – A Typical RF Troubleshooting Process

Courtesy of Doug Kelly, Virgin Media, Ireland

Meanwhile, the cable industry has been investing in remote performance telemetry for a long time, beginning with the DOCSIS 1.0 release in 1997. For nearly 23 years, the DOCSIS specification has been evolving and growing, providing performance metrics for virtually every aspect of our cable networking protocols.

Since DOCSIS 3.0 was launched in 2006, the cable industry has been developing a proactive network maintenance (PNM) specification, which significantly improves the range and depth of available RF diagnostic capabilities. In the hands of a highly skilled technician, these tools can be invaluable and empowering. However, in some cases, they require certain skills and experience, in order to properly interpret and enact a correct repair. When a technician lacks the training or experience to properly use or interpret the tools, they often forego using them and rely a limited repertoire for repair, such as swapping equipment and resetting devices. Our technical workforce is made up of skills on both sides of this continuum and everywhere in between.

As the remote visibility of the customer equipment improves, so do the tools to detect problems, present data and dispatch technicians. Quoting Brady Volpe, owner of The Volpe Firm Inc., “In the mantra of PNM, we are often able to find and repair problems before they impact the performance of the service perceived by our customers.” While that’s an attractive statement to make, it can be difficult to determine exactly which problem is causing the trouble being experienced by the customer. We may be able to proactively detect RF problems before the customer is impacted, but they could be calling due to a completely unrelated problem. In many cases, our tools currently lack the ability to accurately establish cause-and-effect with the customer experience.

© 2020 SCTE•ISBE and NCTA. All rights reserved.

4. Fault Isolation on RF Transmission Lines

To people familiar with customer support and troubleshooting, there are a few relatively simple processes that tend to work universally. Sometimes, problem isolation can follow a basic divide-and-conquer approach, depending on the type of problem. However, RF transmission troubleshooting has special conditions and circumstances that can be problematic for a simplistic troubleshooting process.

It may be helpful to remember that our coaxial cables are essentially just a large, shielded transmission line. Our shielded transmission lines can cover long distances, spanning many splits and taps with little directional isolation, which can befuddle the divide-and-conquer approach. Fortunately, cable has some physical characteristics that help, but do not completely solve these isolation problems. It can be generalized that many problems are difficult or impossible to localize, inside or outside of the customer premises.

The most common and reliable method for automatically localizing problems within a location is the presence of multiple equipment. This allows the remote monitoring systems (and technicians) to compare the signals of multiple sensors within the location. By a process of comparison, all devices within the home sharing a common problem increases the likeliness that the problem is outside (Figure 2). Conversely, if a single device detects a problem while the others do not, it can be concluded that the problem is inside. In recent years, a boom in DOCSIS deployments has made this technique very useful. However, the latest industry trend is moving towards a single DOCSIS point-of-entry gateway, then relying on Wi-Fi to distribute content within the premises. Over time, this will diminish the value of this localization technique.

Figure 2 – Common Problem Signature, Multiple Devices

In a similar process of elimination, when a single problem signature is detected having multiple sensors available (Figure 3), it's a reasonable assertion that the problem is within the location. These types of issues are typical of wiring problems such as loose or damaged connectors, incorrect fittings on wall plates, damaged cables, and splitters that are installed backwards, to name a few. In some cases, the problems can be addressed by the customer without requiring a technician. Some operators use this technique to identify possible loose connectors and inform the customer, instructing them to tighten the

© 2020 SCTE•ISBE and NCTA. All rights reserved.

connector. These types of problems may not cause a noticeable issue for the customer, but tightening loose connectors is one of the easiest and best ways to improve the overall reliability of a coaxial network.

Figure 3 – Single Problem Signature, Multiple Locations

The technique for comparison-based localization improves as more data is introduced, including sensor information from neighboring equipment connected to the same physical network tap (see Figure 4). When common problems are detected that have a shared network element, the shared network element is usually the cause for impairment. These types of problems are commonly found at the tap, and include damaged or corroded tap plates, incorrect termination, cut drops, incorrect pin length, loose pin seizures and more.

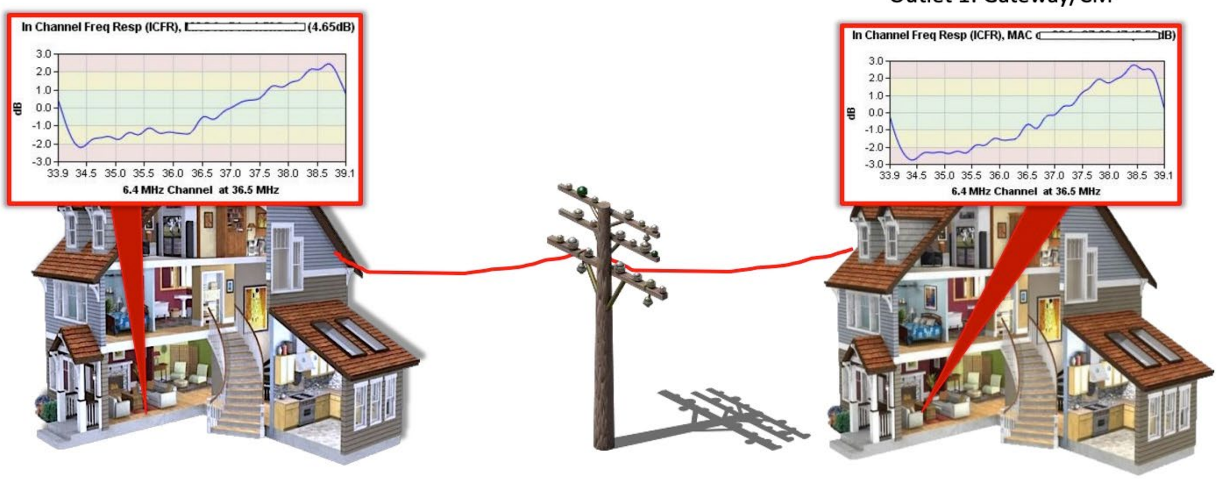

© 2020 SCTE•ISBE and NCTA. All rights reserved.

Figure 4 – Problem Signature Compared to Neighbors

Another useful method for localizing problems uses casual observation and makes inferences based on common signatures. For example, there are certain impairments which are well-known to be outside vs. inside. In the case of water damage, water enters a coaxial cable and creates a distinctive signature that is easily identified (Figure 5). These are nearly always outside, caused by environmental influences such as rain or sprinkler systems. There are possible examples where water can have entered coaxial inside the home, but those are minimal and unlikely. Other examples of inference-based signatures are the presence of filters, which are usually installed at the ground block or tap; old satellite splitters which can produce unique standing waves; and spectral roll-off from old passive equipment, to name a few. These types of signatures are automatically detected with specially designed spectral impairment detection (SID) software libraries available from CableLabs.

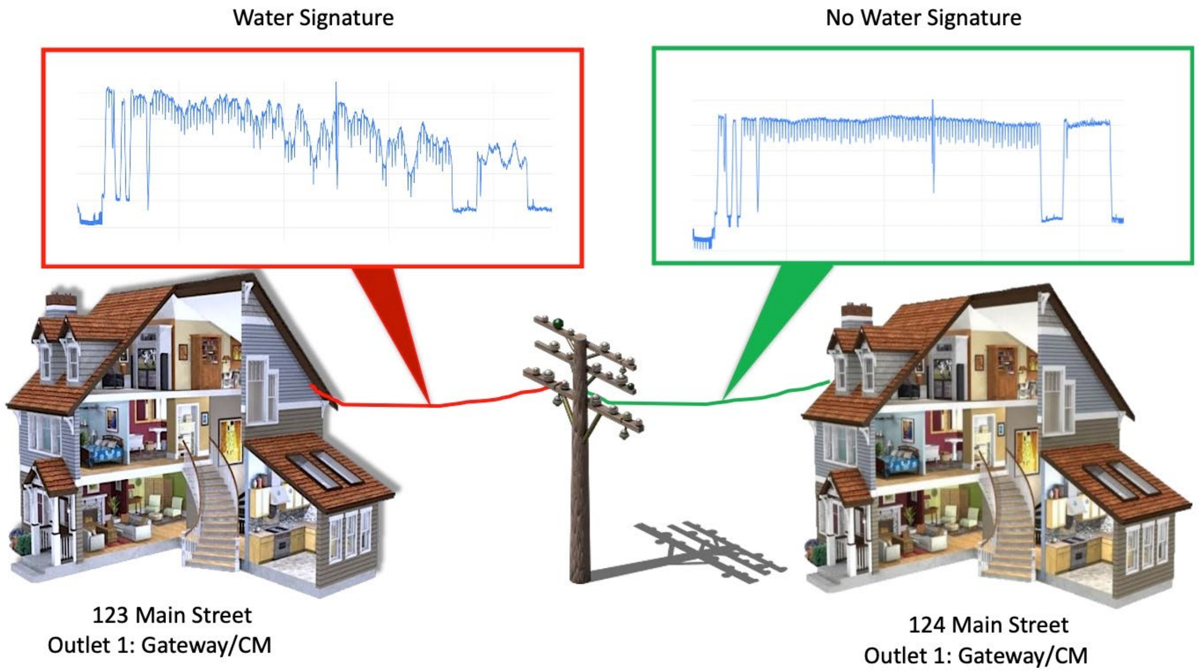

Figure 5 – Localization Based on Signature Inference

The techniques described above work on a subset of the common types of problems that are detectable with our DOCSIS PNM tools. They are especially useful for downstream RF impairments and anything that causes an impedance mismatch, any physical damage on the cable system. As previously discussed, these techniques do not work for certain types of problems. Unfortunately, some of our most impactful and common problems are invisible to this type of isolation technique.

Upstream noise is the most notable exception, including any other spurious types of interference which may be intermittent. The primary reason for upstream noise eluding these techniques is known as the upstream funnel effect. The upstream portion of the RF spectrum is coupled in a manner that allows all of the signals to travel towards the headend, then become combined together at the upstream receiver. The receiver has no way of knowing where this unwanted signal is getting in (ingress), and the noise impacts all cable modems on the same physical RF connection. This is particularly problematic because one

© 2020 SCTE•ISBE and NCTA. All rights reserved.

ingress problem can impact an entire service group, consisting of hundreds of customers. Figure 6 shows the noise present, affecting a large service group. Finding the source of the ingress remains elusive and usually requires a maintenance technician. This is one of the primary motivators to keep our connectors tight.

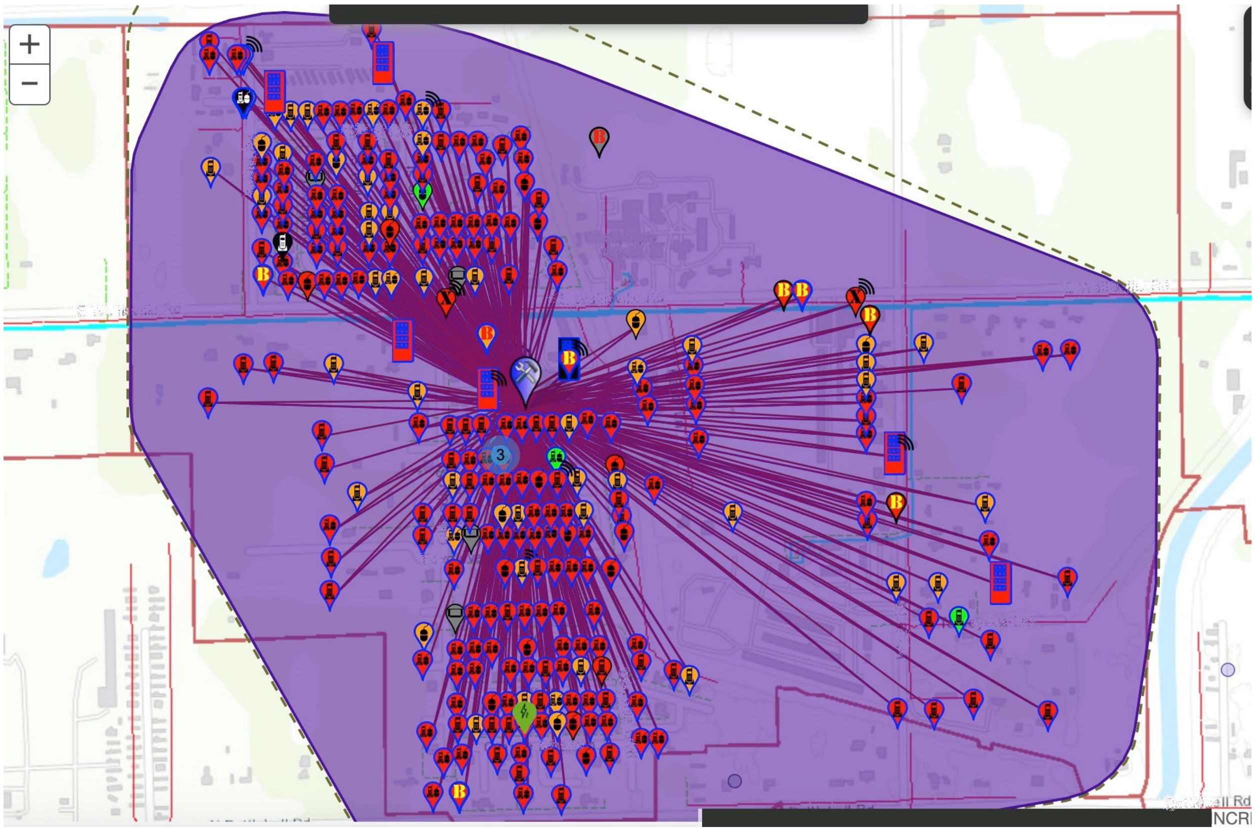

Figure 6 – Upstream Noise Impact on Entire Service Group

Occasionally, some upstream noise problems are localizable to a small area, but these are relatively uncommon compared to the more typical upstream ingress. These types of problems can be caused by faulty return amplifiers which have a detectable distortion effect on the upstream signal for a smaller group of cable modems. With these "pocket issues," it sometimes works to group devices that have a common RxMER or FEC-related problem. Unfortunately, this technique is minimally effective at localizing most typical upstream noise or ingress problems.

4.1. The Troubleshooting Process

In days prior to COVID-19, the troubleshooting process typically would begin with the technician making contact with the customer to obtain their perspective, in-person. This can transpire any number of ways and usually relies on customers using their own language to describe the service problem. For this to be effective, the technician needs to interpret the customer's description and translate it to any number of troubleshooting and repair processes. For example, the customer may describe the problem as "the cable isn't working," which would require additional questions to further inform a) which cable service (video, voice, broadband) and b) which troubleshooting process to use. Next, the technician may ascertain that the customer was referring to internet service, then further determine if the connection is Wi-Fi or wired, and so on.

© 2020 SCTE•ISBE and NCTA. All rights reserved.

This is the input-side of the instruction loop, usually referred to as "features" in machine learning nomenclature. It's easy to understand how translating the customer's intent can result in misunderstood features. David Monnerat, Sr. Director of AI at Comcast, noted, "I've been on truck rolls in three states, all good people trying to do the right thing. They have good knowledge and skills but not a consistent process, all troubleshooting differently. Three out of five technicians would have a different resolution to the same problem." We're faced with a wide variety of products, experience and technical understanding across our customers, agents and technicians. Fortunately, our tools and technology have been evolving to help bridge these divides and produce more consistent outcomes for our customers.

4.2. Repair Feedback

After completing repairs, the technicians are usually required to provide some form of feedback about the work that was done to resolve the customer-reported problem. Typically, this is done when a technician closes a job (work order), and they are prompted to provide some mandatory selection from a list. This list can be vast and may allow for multiple selections, which will be discussed further. Most operators have used systems similar to this and experienced the same results regarding the unintentional bias in the resulting feedback. For instance, technicians might arbitrarily select the first code from a long list, regardless of the code.

4.2.1. Technician Performance Metrics

Among of the most influential drivers of bias within the repair-feedback loop are the performance metrics used to measure technician productivity and effectiveness. These metrics are used for reports that assess how well technicians are doing, ultimately resulting in pay/career growth or the opposite, relative to their performance against the operator's established goals.

Lessening repeat or re-work is a common goal when a service call requires multiple technician visits within a specified time period, such as a month. Given the many possible scenarios of service repair, there are loopholes which can insulate technicians against demerits from having to go back at a later time. Depending on the reporting algorithms, a technician may improperly report that the customer was not home, even though they may have performed some repair activity.

Another of the more common re-work loopholes is to complete the work order as avoidable (not required) by coding the repair with something like "no trouble found." In this case, a technician may perform some repair activity, but depart unsure that it actually solved the problem. If they think there is a reasonable chance that the problem may be intermittent and unsolved, this provides some cover for a repeat visit in the future.

Although it doesn't directly impact the repair coding of work orders, it's not uncommon for technicians to leave their personal contact information with customers, so as to call them directly. This is another way for them to circumvent demerits associated with re-work. If the technician needs to go back and perform additional repairs or support, they do it off-record, avoiding negative performance reporting.

4.2.2. Customer Relationships

The financial policies of cable operators can often influence technician repair feedback. It is not uncommon for operators to have a policy that establishes rules about when customers should pay for the service call versus it being free-of-charge, at the expense of the operator. These service call fees are often between \$50 and \$75, depending on the circumstances. Examples of chargeable service calls include scenarios such as a customer incorrectly re-connecting equipment when rearranging furniture, or if the power was simply turned at a power strip. In either of these examples, a minor oversight by the customer

© 2020 SCTE•ISBE and NCTA. All rights reserved.

could prove embarrassing and is unfortunately all too common. Many technicians can empathize in these situations and find some menial form of repair that can be done to avoid an uncomfortable conversation about having to charge the customer for their oversight. It's simple enough to code the work order as replacing a connector or splitter, versus having to break the bad news about a charge for the service call. There are many areas where technicians have never charged a customer for a service call, providing anecdotal evidence that this is an all-too-common scenario.

4.2.3. The Expert Technician / Opportunistic Repairs

There is a population of technicians that is especially enthusiastic about performing high-quality work. These technicians take pride in leaving every customer in better condition than before they arrived. They will scour the premises, starting at the tap, looking for loose or corroded connectors, then inspect the drop for leaks or damage, continuing to the ground block for distress or connector problems. These experts will invariably find plenty of opportunities for proactive repair, such as replacing old F-connectors, corroded ground blocks and imperfect cables. These are all great qualities that we hope all of our technicians would exhibit. However, they will often reflect every aspect of the repair in the work order. The multitude of proactive / opportunistic repairs may not have affected the original reason for the trouble call, although they certainly represent great hygiene for the network and insulation against a repeat trouble call. Also, by including multiple repair codes, this sometimes influences their performance statistics in a positive way. Naturally, this can result in an inflated number of repair codes that might not relate to the customer-reported problem.

4.2.4. Report Burden and Selection Bias

Although it may not seem over burdensome to provide thorough feedback after a service call, the reporting process after a long, hard repair job can prove mentally exhausting – especially if the technician doesn't believe the outcome of the report will have a meaningful effect on them or the business. In business terms, this represents a form of "decision fatigue," resulting in the technician not putting sufficient energy towards making a proper selection to describe the job. This is exacerbated with growing lists, of hundreds of codes, which can take a long time to scroll, read and contemplate. Several controlled studies at Comcast have shown that the top code arbitrarily gets picked the most. Further attempts to randomize the top selection result in randomization of the repair disposition. The results were predictably consistent with the first presented code.

Some operators have attempted to reduce the selection burden by ordering the codes by their predominance. For example, replacing equipment, resetting devices and reprovisioning service are usually among the top selected codes. As a matter of improving the user experience, these codes might be ordered as the top three in the selection list, to reduce the searching and scrolling required by the technicians. Unfortunately, by placing the top selected code in first place, the previously discussed problem of arbitrary selection becomes compounded. This creates a selection loop bias which further strengthens the predominance of a small number of repair codes used to characterize a repair.

4.2.5. Intermittent Problems and Upstream Noise

The nature of RF performance and troubleshooting can sometimes be intermittent. Upstream RF noise is notoriously spurious in nature. The source of the noise may be intermittent, such as turning on electrically noisy equipment, like the electric motors used in hair dryers or power tools. It's also possible that the place the noise is getting in may be intermittent, such as a loose connection on a drop cable that might be blowing in the wind. As the wind blows, this can cause unpredictable shielding faults within the connector's threads that intermittently allow the noise to enter the drop. If a customer calls at the time of the ingress and impaired service, it may be hours or days until a service technician may be able to visit the

© 2020 SCTE•ISBE and NCTA. All rights reserved.

customer. It's also possible that the ingress may be coming in at a completely different location on the node, unrelated to the customer who's calling about the problem. In these cases, technicians will typically attempt some hygienic repairs, such as replacing connectors, without being able to confidently assert that the underlying problem was fixed. In these examples, the technician may code the repair as no trouble found, replaced connectors, swapped equipment or any number of things.

4.2.6. Insufficient Tools or Training

As our services become increasingly more complex, it can take time for our technicians to learn how to properly troubleshoot and repair them. For example, when Comcast introduced DOCSIS 3.1 service, there was a learning curve and tooling upgrades that were required to diagnose and repair orthogonal frequency division multiplexing (OFDM) signals. There was a time when legitimate RF problems could have been impacting a customer's experience, but the technician lacked proper tools or training to identify and repair the issue. There are other problems such as capacity, congestion, software updates and a myriad of others that may not be presented to the technician's troubleshooting process. The result of an undiagnosed problem typically results in the "Hail Mary" approach. Left with no other options, a technician will often fall back to replacing the customer equipment, otherwise known as a box swap. The vast majority of times, the equipment is not at fault, but it is possible that this activity does improve the service. At the very least, it offers the technician an opportunity to demonstrate that they are doing something perceived as helpful. In some cases, the act of swapping equipment usually includes re-provisioning, which can help. An example would be correcting an incorrectly-provisioned boot file which was unnoticed by the tools or technician. It could also be possible that a new device would provision to a different RF channel set which is less impaired than the previous implementation.

One of the most common forms of service repair is simply resetting or rebooting the equipment. While this does nothing to physically repair RF problems, it can sometimes be useful and has a very low cost, other than temporarily disrupting the service. In cases where the equipment software may be having problems, rebooting the device can temporarily re-initialize the software and restore proper function. However, this tends to be a temporary fix until the software bug is encountered again.

The latest versions of DOCSIS have proven to be exceptionally resilient. Our DOCSIS specifications have many coping mechanisms available to enable operation even in the most hostile RF environments. For example, a DOCSIS 3.1 cable modem may have 32 or more downstream channels (in addition to the OFDM signal) available for use. It is not uncommon to see frequency-specific problems, such as LTE ingress, which can impair a few channels, while the others might be operating perfectly. DOCSIS can bond different channel sets or disable problematic interfaces to allow error-free operation. This is another example of how modem resets can restore service without performing a repair to the physical environment.

5. Cloud Connected Technicians and Equipment

Our field signal meters, and measurement systems have also been evolving, becoming more connected and integrated with the technician ecosystem. In our contemporary workforce, the signal analyzers are cloud-connected and augmented with all kinds of new information, such as technician identity, GPS coordinates, system design maps and telemetry measurements from customer equipment. This information-rich environment creates new opportunities for features and labels to provide to our ML and AI systems.

© 2020 SCTE•ISBE and NCTA. All rights reserved.

5.1. A New Troubleshooting Process

As discussed previously regarding the troubleshooting process, it becomes obvious that this is an area of opportunity for improvement. With the onset of COVID-19, our customers, technicians, agents and business partners now have an entirely new set of constraints and motivators shaping how they approach the troubleshooting process.

The first and most obvious constraint is the desire to maintain physical isolation between our technicians and customers when servicing their equipment. Brady Volpe, when asked about his experience on RF related repair calls, cites that "about 75% of the time it's tap, drop or ground block. The other 25% of the time is in-home wiring … these statistics exclude many of the common trouble call issues, such as customer education, and Wi-Fi problems, and varies by area." Going by those statistics, and as a matter of efficiency, it makes sense to start the troubleshooting process outside. At Comcast, a new process was devised to help seize upon this statistical advantage and provide additional safety by allowing for physical separation. In addition to those two key benefits, it also facilitates a new opportunity for a consistent troubleshooting process.

As seen in the flow chart (Figure 7), technicians now start all RF troubleshooting by taking a signal measurement outside of the service location at the ground block. This is a critical demarcation, indicating where the cable service becomes physically attached and electrically bonded to a service location.

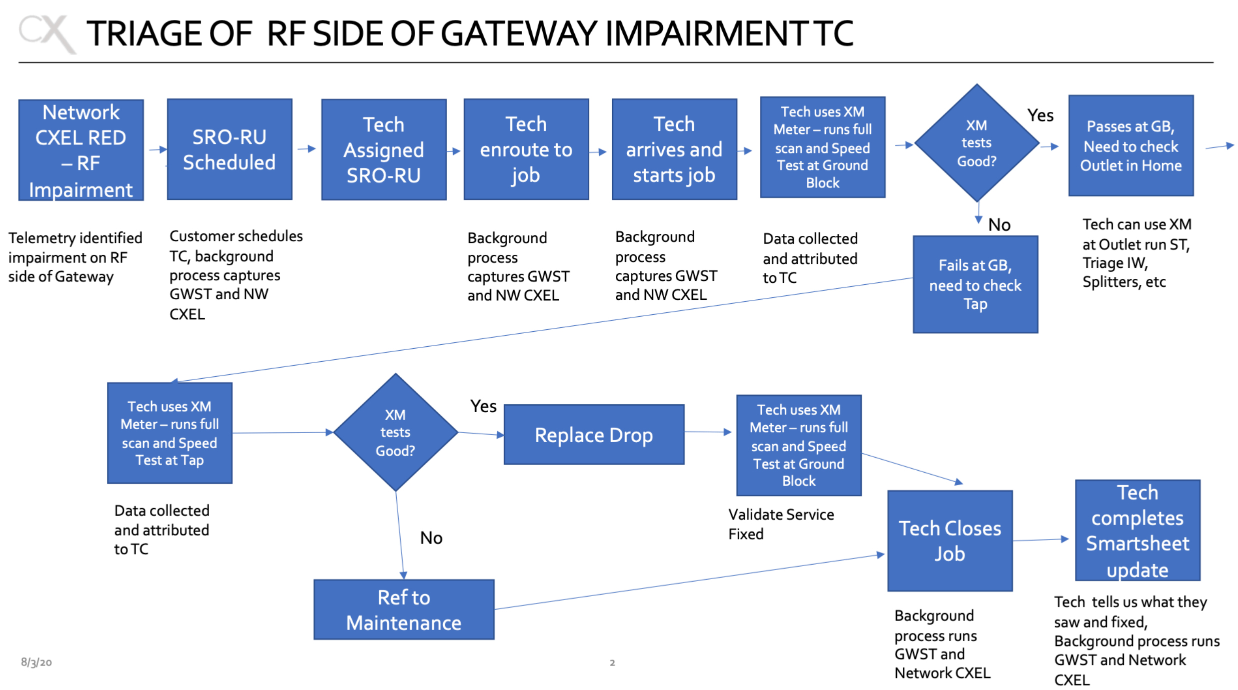

Figure 7 – New Troubleshooting Process

5.2. Process Compliance

Shortly after implementing cloud-based signal measurements, in early 2017, issues of process compliance became obvious. This was one of the first and most intriguing insights being provided by our new field

© 2020 SCTE•ISBE and NCTA. All rights reserved.

measurement platform. A meager 22 percent of technicians were using their meters in a 24-hour period on any repair jobs. This quickly leads to all kinds of other questions about the troubleshooting process. How can someone possibly diagnose, repair and validate service without taking measurements? It stands to reason that not all trouble calls require RF troubleshooting; however, these 22 percent of repair calls were service-related and resulted in RF specific repair codes.

At Comcast and a number of other operators, it has become common place to use the premises-located equipment to measure signal levels and performance. One of the leading causes for process non-compliance is that the technicians were relying on a premises health test (PHT) in lieu of taking measurements. This process evaluates the remote telemetry from the DOCSIS gateway and other equipment. While this is an acceptable way of validating service conditions, it's a procedural shortcut for the troubleshooting process. Lacking portability and segmentation of the network, this method results in a significant amount of "hunt and peck" rather than accurately diagnosing and repairing problems.

6. Machine Learning and Artificial Intelligence

ML and AI have become ubiquitous in modern computing systems. Cable operators are investing heavily in these platforms and continue to find opportunities to improve how we operate our systems. This section reviews some of the techniques used, and results achieved when attempting to segment the RF network using ML and AI.

6.1. Features and Labels

The review of the troubleshooting process was important to help convey an understanding of the features that we'll be using for our models. Features and labels are important constructs in ML and AI for both classification and regression problems. Features can be thought of as the inputs that will be available to the model. The features, in our case, will be any number of the measurement data which are available to our systems. Features might include in-channel frequency response (ICFR), full band capture (FBC) impairments, packet loss, speed test results, time of day and number of device resets.

Labels can be thought of as the output, or the desired prediction from the model. Examples of labels might be loose connectors, faulty drops or excessive splitters. These are the outcomes that we would expect our models to predict, given the proper features to inform its decision tree. Notice that the labels mentioned do not include customer behaviors, such as trouble calls or other service requests and interactions. One important aspect of this exercise is to decouple the objective service conditions from subjective customer experience. For instance, poor RxMER and packet loss may result in insufficient speed test performance. Some customers may call while others may not. That is discussed later.

© 2020 SCTE•ISBE and NCTA. All rights reserved.

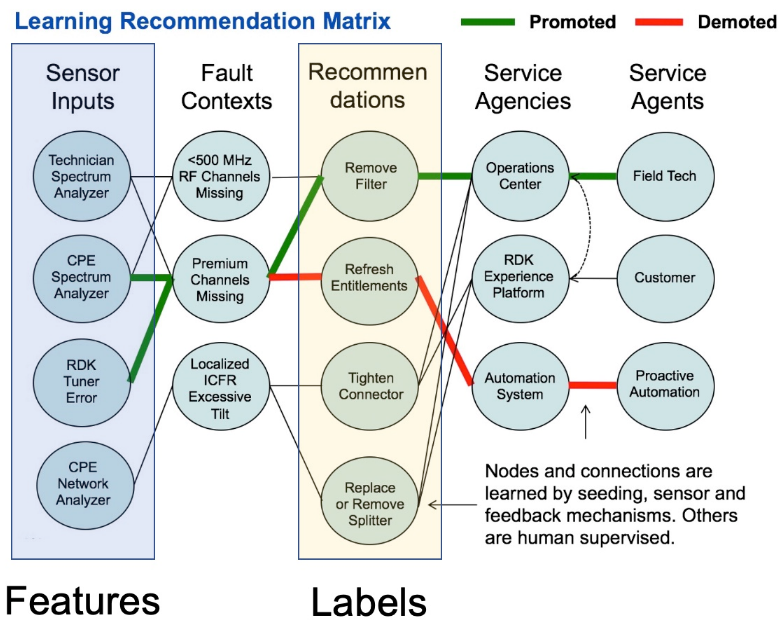

Figure 8 – Example of Cable-Oriented Features and Labels

6.1.1. Machine Data vs Human Data

The previously discussed, fallibility of the human interpretation in our feedback system can become one of the most promising elements to improve our ML and AI model training potential. "You can think of machine learning using familiar terms to cable engineering. The goal of our models is to pick out the signal from the noise," stated Jan Neuman, Executive Director of Machine Learning at Comcast. "Objective measurement that is noise-free and repeatable can be used to refine other measurements. By adding a less noisy signal, this increases the overall fidelity of the model. Thus, increasing the signal-to-noise ratio (SNR) results in more accurate predictions." Put another way, improving prediction accuracy is the primary motivation for removing noise from our models.

Looking back on the discussion about the troubleshooting process, the first opportunity for noise begins at the input, or features. Although difficult to quantify empirically, the most fundamental features – such as "why did the customer call for help?", or "which service is having a problem?" can be corrupt. If customer language is interpreted literally, "the cable isn't working" could have many different meanings, causing noise in the feature. When this noisy feature is incorporated by the model and it starts making predictions, it is difficult to imagine that the results will be useful. By using machine data to refine the feature definition, the noise is reduced at the input of the model. An example of improving this feature with machine data would be a technician testing RF the ground block, outside of the location. If significant packet loss is measured, a noise-free feature now exists that is causal for a poor internet experience.

In addition to refining the features, machines offer an opportunity to improve the fidelity of the labels, or desired outcomes and predictions. While still not perfect, there are machine-provided data that can be reasonable proxies to some of the common labels. For example, the label of slow internet speed can be

© 2020 SCTE•ISBE and NCTA. All rights reserved.

approximated with automated speed testing. As expected, features of packet loss and labels of slow speed are clear machine signals that can be interpreted without added human noise.

6.1.2. Digital Interactions With our Customers

For decades, our customers have been placing telephone calls to our call centers. As we drive towards more digital interactions with our customers, we're offered a unique and important opportunity to remove some degree of human interpretation (noise) from the models. Cable operators are embracing digital interfaces for our customers such as apps and online tools. By allowing customers to directly convey their intent, this offers an opportunity to bypass one of the most common causes of noisy labels, human-to-human communication on the telephone.

6.2. The Model

Comcast Applied AI researchers developed a classification model to identify service calls where only outside of the premises work was required. The team developed and deployed this model to schedule trucks with a goal to minimize contact between technicians and customers during the COVID-19 isolation policy. After the initial model was deployed, they explored additional data, such as field meter measurements, to improve the model's performance.

6.2.1. Outside of Home Applied AI Model

The "Applied AI outside of home" machine learning model used features from several network telemetry sources that collect and aggregate DOCSIS measurements. These sources poll and analyze the network, looking for outages and impairments. Other systems collect network data and do the fault segmentation analysis described earlier, which was also included. Table 1 describes the different data sources.

Source	Description
Account type	Account-level detail showing types of devices in the home (video gateways, wireless gateways, modems, etc.), days since account origination and since last device installation.
Prior truck history	Summary of technician-reported problems aggregated at the node level.
Account network degradation algorithm	Features developed from an algorithmic tool that polls devices four to six times per day to report account-level degradation issues related to disturbances in the RF spectrum. Features are based on issue counts related to network, drop, in-home wiring, loose connections, and isolated home concerns.
Connectivity between modem and CMTS	Features developed from a tool that polls devices three times per day to report raw measurements describing connectivity between the modem and the CMTS. Features include counts of impairment flags and means

© 2020 SCTE•ISBE and NCTA. All rights reserved.

	and standard deviations of measurements. Data reported include timeouts, system boot time, ripples, SNR, transmit and receive power, FEC errors, tap energy, and phase angle deviations.
Node network impairment algorithm	Features developed from an algorithmic tool that scans the network for RF impairments and reports them continuously. Impairments include plant Wi-Fi, suckout, wave, flux issues.

Table 1 – Model Sources Summary

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 down to customers' homes (Figure 9). 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 a weighted average over a customer's node, where the weight is determined by the graph distance from the customer with the truck roll to other customers in the node. Another type of aggregation is at the parent level, which averages the telemetry over all customers who are immediate neighbors in the graph. We also computed averages for each piece of network equipment and compared it to averages for neighboring equipment, then found the piece of equipment that the customer depends on which has the greatest (and least) difference from its siblings. The idea is that if a single piece of network equipment is broken, we can observe 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.

© 2020 SCTE•ISBE and NCTA. All rights reserved.

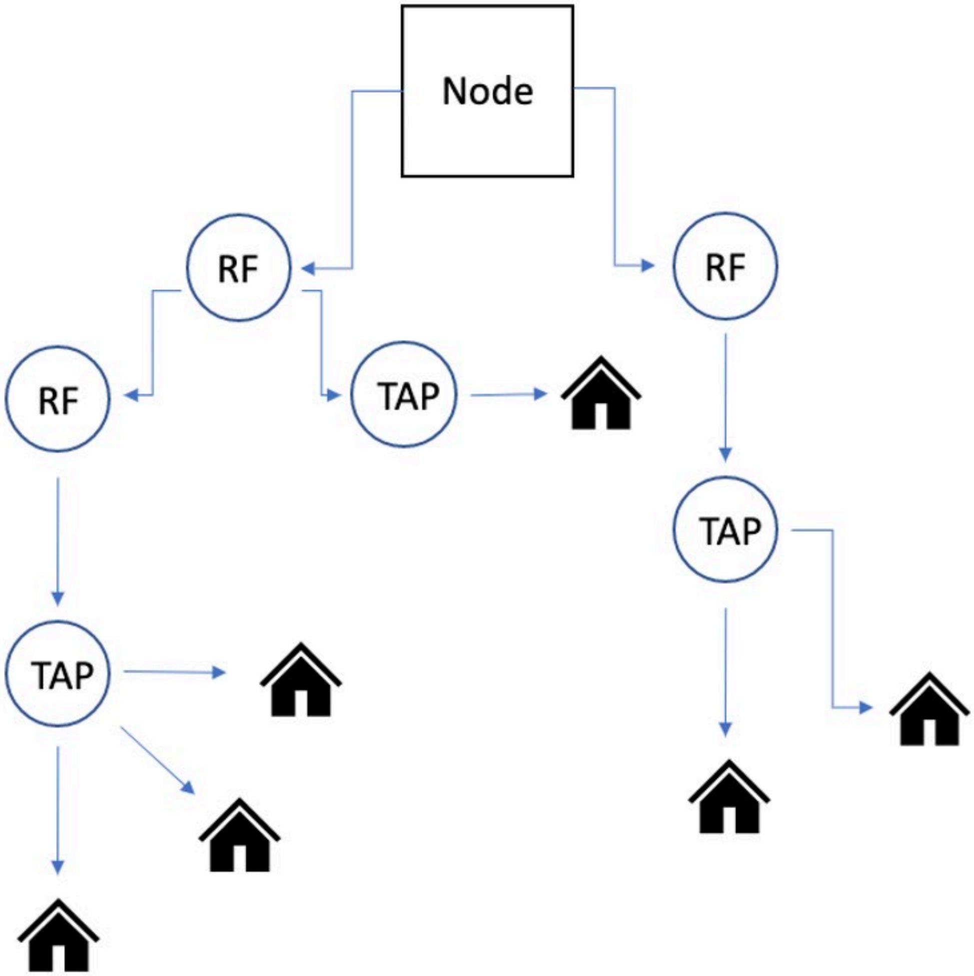

Figure 9 – Node Topology Schema

Finally, the feature engineering pipeline aggregated impairments and telemetry data in a daily timeframe, predicting whether the truck roll would only require outside work.

The model supplied predictions to the agents, who advised customers on the type of truck to schedule. The team defined the target label using repair codes, provided by technicians after the service was completed. The outside fix codes included “refer to maintenance,” “construction,” or, for underground teams, “replace connector,” or “replace, repair, or run underground drop.” The best performing model was trained with an-open source XGBoost classifier. The model was calibrated to achieve at least 5% recall and reported precision at 77% and lift at 2.48. Figure 10 illustrates the precision versus recall curve, normalized to a value of 1. Figure 11 shows the individual feature importance.

© 2020 SCTE•ISBE and NCTA. All rights reserved.

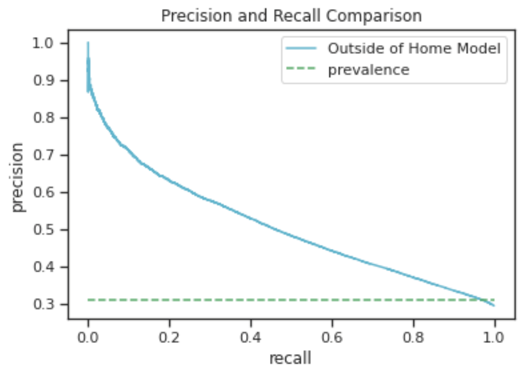

Figure 10 – Model Performance, Precision and Recall

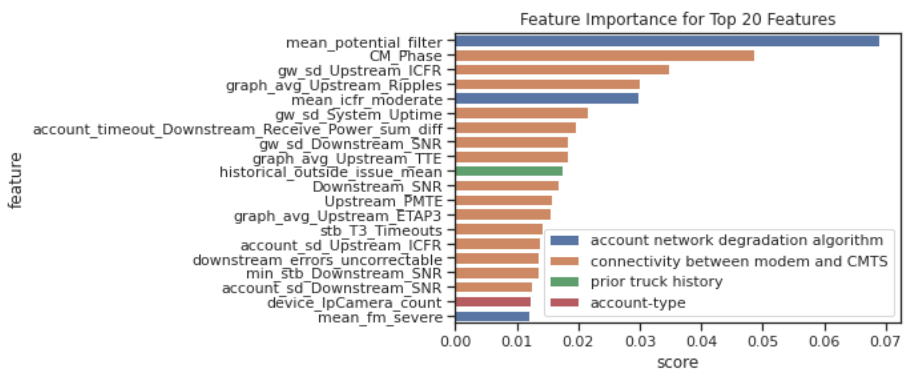

Figure 11 — Model Feature Importance

6.2.2. Using Weak Supervision to Improve the Model

After the model was deployed, the operations team launched a survey asking technicians to specify whether work was needed only outside of the house. These data were considered to be ground truth to verify the output predictions. Using these labels could potentially improve the model performance. However, ground-truth labels data were small and modern machine learning techniques require a large volume of data to train the model.

To address this problem, the team applied weak supervision techniques using the open-source Python package Snorkel (https://www.snorkel.org/). The goal was to access data not available at the time of prediction, use those data to develop a labeling model, train the labeling model on the ground truth data, and assign labels to the larger corpus of data that was lacking ground truth labels. Then, our data size would be sufficient to train a classification model on labels close to the ground truth and with sufficient data size. In other words, we attempted to use the results of signal meter readings after a small number of repairs, to predict what will happen on all trouble calls before the technician gets there. Figure 12 illustrates the weak supervision approach.

© 2020 SCTE•ISBE and NCTA. All rights reserved.

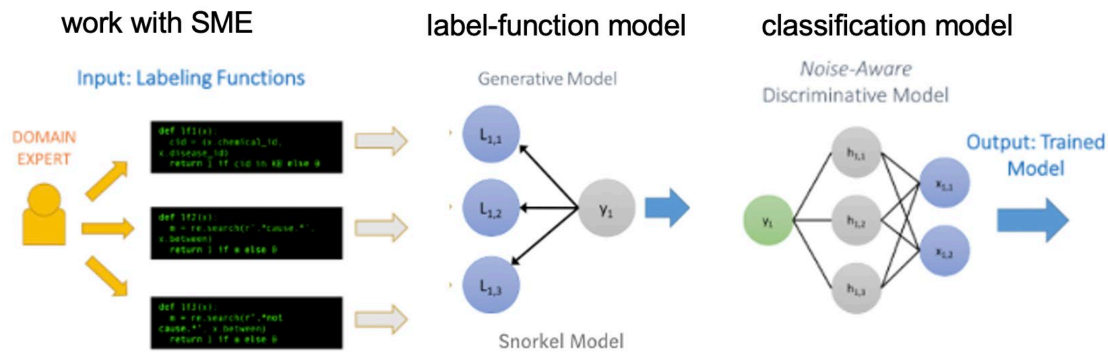

Figure 12 – Weak Supervision Workflow

The team gathered a corpus of data to use specifically for label development. Many of those sources were not available at the time of prediction. Those data included technician comments, repair codes, and signal meter data, among others.

Technicians took signal measurements on-site at the time of the truck visit. Of the total number of repairs, approximately 50% of them had field meter measurement data available. Engineers pre-processed the data for the team and developed labels by aggregating measurements such as per-channel RxMER, receive power and ingress. They assigned the following labels to each measurement: “ground block pass,” “ground block fail,” “tap pass,” “tap fail,” “refer to maintenance,” “tap,” and “tap fail ground block pass.” Accuracy of the functions derived from signal meter data ranged from 55% to 62%, surpassing the prevalence of outside labels in the survey data (at 45%). Table 2 enumerates the function performance.

© 2020 SCTE•ISBE and NCTA. All rights reserved.

function	case	coverage	overlaps	conflicts	accuracy
lf_survey_outsidefittings	outside	5%	5%	4%	100%
lf_problem_tap	outside	11%	11%	7%	100%
lf_survey_abletorepair	outside	33%	33%	24%	100%
lf_problem_groundblock	outside	5%	5%	4%	100%
lf_problem_drop	outside	18%	18%	13%	88%
lf_problem_refer_to_maintenance	outside	7%	7%	5%	82%
lf_problem_refer_to_underground	outside	4%	4%	3%	80%
lf_xm_overallresult_gbfail	outside	6%	6%	3%	62%
lf_xm_overallresult_drop	outside	2%	2%	1%	61%
lf_xm_outsidelabel	outside	16%	16%	9%	59%
lf_xm_overallresult_gbpass	inside	16%	16%	16%	59%
lf_xm_overallresult_rtm	outside	4%	4%	2%	57%
lf_xm_overallresult_tappass	inside	12%	12%	12%	56%
lf_xm_insidelabel	inside	41%	41%	41%	55%
lf_xm_overallresult_tapfailgbpass	outside	3%	3%	2%	55%
lf_survey_issuefixed	outside	54%	53%	40%	55%

Table 2 – Labeling Functions with Highest Empirical Accuracy

An important aspect of this exercise was the collaboration with data and field engineers to write functions and rules about the ground truth label. For example, if meter test interpretation was “ground block pass” or “tap pass,” the label model assigned the “inside of home” label. If test interpretation resulted in “drop,” for example, functions offered the “outside of home” label. Each function conveyed a proposed rule for labeling “inside of home” or “outside of home” case.

Many of the functions used technician-entered repair codes and free-form text comments. A set of these codes detecting tap and ground block issues had very high accuracy. Other repair codes varied in accuracy but made great contributions to the labeling model when they were combined with functions derived from signal measurements.

© 2020 SCTE•ISBE and NCTA. All rights reserved.

42 total functions were developed based on the combined labeling data corpus. Naturally, resulting functions varied in accuracy and coverage; some had labeling conflicts. We pruned the functions to include only those with higher accuracy. Then, we applied a majority voter strategy to assign predicted ground-truth labels and resolve function conflicts. In this strategy, cases where functions offered contradicting predictions were labeled neither inside nor outside in the training data.

Finally, a classification model was trained using Snorkel-labeled data. Both, benchmark and Snorkel-labeled models were trained on equal feature sets using the XGBoost classifier. We retrained both models on a limited set of data to benchmark the gains that can be achieved with the weak supervision method. Table 3 shows the precision at 5%+ recall target. The model trained with the snorkel-labeled data exceeded the performance of our benchmark model, with precision increase from 62% to 68%. By incorporating signal meter readings collected on-site, we were able to improve the accuracy of our labels and our model.

Model	**Precision**	**AUC**	**Prevalence**	**Lift**	**Data size**
Ground truth label (benchmark)	62.15%	0.545	45.8%	1.35	34,860
Snorkel-labeled	68.32%	.555	45.8%	1.49	370,154

Table 3 – Model Performance Comparison

7. Conclusion

Although all manner of misdirection (noise) exists within the data, there are still valuable insights to be gained. Our research has demonstrated that weak supervision techniques, access to subject matter experts, and a corpus of data are useful in developing labeling functions can help us improve the model performance. Specifically, access to diagnostic tools such field meter measurements can help us to improve training data labels.

Machine learning is a logical next-step toward identification and isolation of problems in the RF plant. Adding new data – in this case, our signal meter data – enhances RF domain expertise, improves operational performance and eliminates repeat service calls. By converging the ML/AI predictions with a real-time recommendation and feedback loop, there are also operational improvements to be realized (Figure 13). In other words, by using the results learned from highly effective technicians, the entire workforce can improve, resulting in more efficient technicians and improved customer experience.

© 2020 SCTE•ISBE and NCTA. All rights reserved.

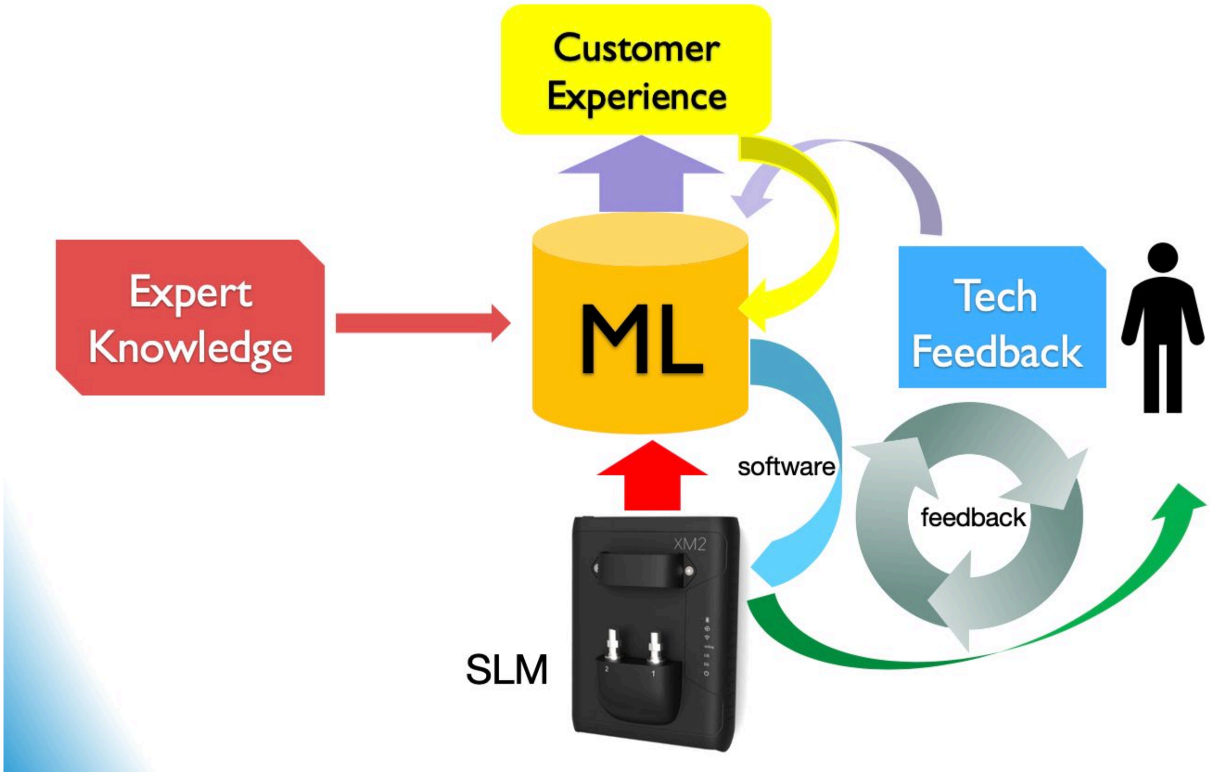

Figure 13 – ML/AI Enhanced Process and Feedback Loop

Abbreviations

AI	artificial intelligence
CMTS	cable modem termination system
DOCSIS	Data-Over-Cable Service Interface Specifications
FBC	full band capture
FEC	forward error correction
GPS	global positioning system
ICFR	in-channel frequency response
LTE	long term evolution
ML	machine learning
OFDM	orthogonal frequency division multiplexing
PHT	premises health test
PNM	proactive network maintenance
RF	radio frequency
RxMER	receive modulation error ratio
SID	spectral impairment detection
SLM	signal level meter
SNR	signal-to-noise ratio

© 2020 SCTE•ISBE and NCTA. All rights reserved.

Bibliography & References

[1] CableLabs DOCSIS® Best Practices and Guidelines - Proactive Network Maintenance Using Pre-equalization. 2012

[2] SCTE Cable-Tec Expo 2016 - A Comprehensive Case Study of Proactive Network Maintenance; Wolcott, et al

[3] Snorkel - A Weak Supervision System, May 31, 2019; Shreya Ghelani

© 2020 SCTE•ISBE and NCTA. All rights reserved.