How AI/ML drives the evolution toward autonomous broadband networks

White Paper

This white paper examines how Artificial Intelligence (AI) and Machine Learning (ML) are driving the evolution of autonomous broadband networks. It explores a “sense, think, act” framework that enables more data-driven decision-making. Key topics include anomaly detection, digital twin networks, and closed-loop automation for intent fulfillment and assurance. The paper also discusses how AI/ML enhances operational efficiency, enables predictive maintenance, and improves human-machine interaction in network management. By showcasing current applications and future possibilities, it demonstrates AI/ML’s transformative impact on broadband networks, positioning them to meet evolving digital communication demands.

Contents

Introduction

In the 21st century, digital communication technology has become an integral part of society, reshaping the way we live, work, and interact. As this digitalization pervades our daily life, we depend increasingly on broadband. For households, a reliable high-speed internet connection is essential for modern life, while for enterprises, it is vital for operational efficiency and competitiveness. Therefore, the broadband market has become more challenging: the variety of services and applications has grown tremendously, and consumers are more aware of their needs.

The growing complexity of network management necessitates a higher level of automation to streamline operations, minimize human error, enhance productivity, and optimize service quality. Modern software-defined access networks introduce data-driven decision-making and closed-loop automation. The data-driven decision-making is enabled by innovative push-based streaming telemetry and always-on network configuration in the cloud. Closed loop automation is essential to intent-based networking (IBN) and the evolution to autonomous networks. IBN enables networks to translate high-level intents into actionable network configurations and to continuously adjust to maintain desired outcomes effectively. More information is available in the application note "Software-defined access networks" and white paper "Broadband network telemetry". The Internet Research Task Force defines IBN in RFC 9315 while TM Forum offers a technical architecture for autonomous networks in IG1230.

Networks that sense, think and act

Artificial intelligence (AI) and machine learning (ML) are the foundation for autonomous networks and data-driven decision-making. AI is defined as any human-created device or system capable of perceiving its environment (sense), making decisions (think), and taking actions (act) to maximize its chances of achieving a goal. Within AI, machine learning serves as a sub-category that enables computational systems to learn tasks without explicit programming. ML achieves this by recognizing patterns in data and autonomously learning useful features from it.

AI/ML-based decision-making can be supported by a digital twin network (DTN): a virtual representation of a physical network. Such a digital replica is created with direct linkage to operation data from the physical network and is synchronized during the entire life cycle. Because the DTN reflects the real-time state and behavior of the physical network, it can be utilized for analyzing, diagnosing, and controlling the physical counterpart.

A network domain controller implements the "sense-think-act" framework (see Figure 1), which is a conceptual model used to describe the fundamental processes involved in self-managed autonomous networks. The addition of AI/ML-driven control enables dynamic traffic management, data-driven capacity planning, predictive maintenance, and automatic fault resolution.

Intent fulfillment

The first step in operating an autonomous network is for the operator to ingest intent (outer control loop 1), which means that the intent is obtained through interactions with the operator. The human-machine interaction can be made easy and natural by using large language models (LLM) making intent ingestion accessible to a wider range of users beyond expert network engineers. Additionally, intent can even be learned automatically, for example, from traffic patterns or the device used (sense).

Suppose, for example, that the operator aims to set up a multi-megabit internet service for a specific user. The controller translates this abstract intent into concrete configuration. Before applying the configuration to the network (act), the controller may use the digital twin network to validate the feasibility of the intent (think). This feasibility depends on the network capability and usage. If the intent is not feasible, the DTN can suggest an alternative intent.

Figure 1. Sense-think-act framework implemented in a network domain controller. The inner control loop (orange) is fast-paced and fully automated. The two outer control loops extend to the user space and may require human intervention.

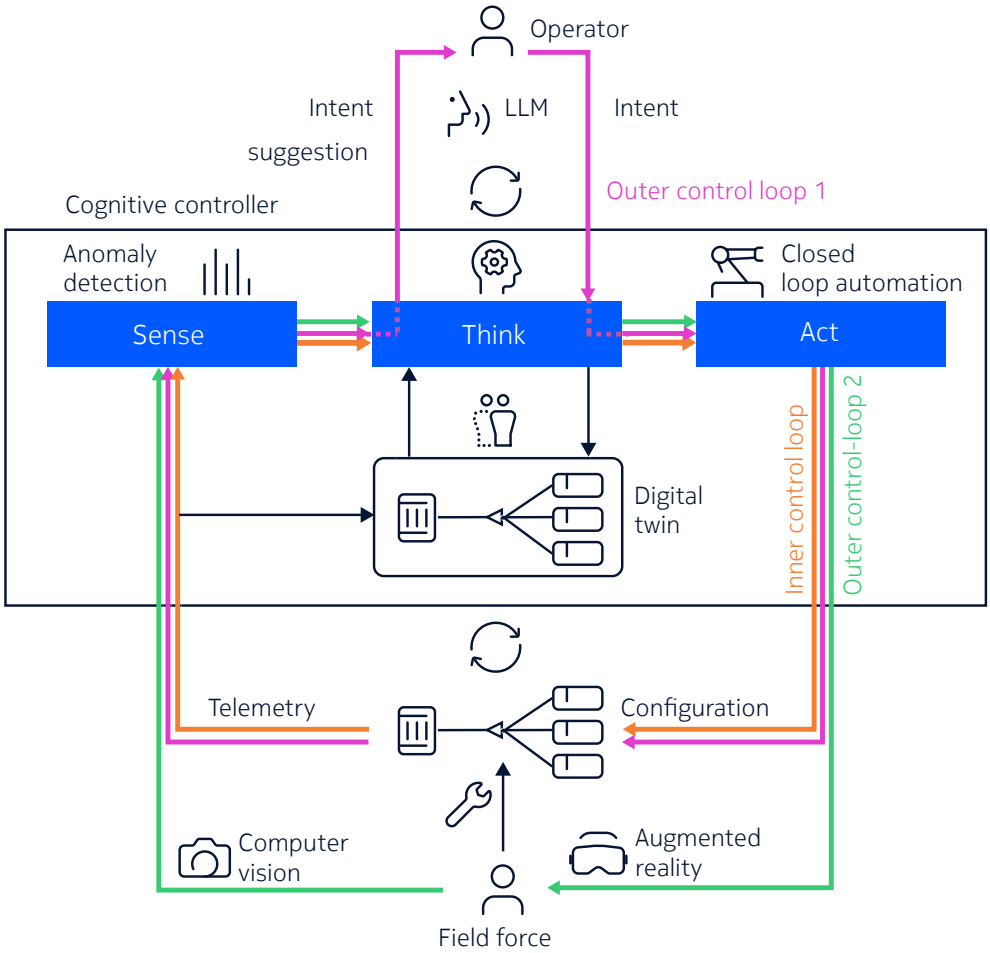

Intent assurance

Once the service is configured, the network monitors the service (inner control loop and outer control loop 1). The access node continuously streams network data to the controller using push-based telemetry. Upon data collection, the controller starts the data analysis using anomaly detection and pattern recognition techniques (sense). If a problem is identified and prioritized, the controller determines autonomously the appropriate corrective action (think). The decision-making may be supported by the digital twin network. The corrective action may be an intent suggestion validated by the DTN (outer control loop 1), a network reconfiguration (inner control loop), or a repair advice to be handled by a field technician

(outer control loop 2). If the corrective action requires an adaptation of the network configuration, such a reconfiguration can immediately and automatically be applied to the network (act). Therefore, the inner control loop is fast-paced and implements closed-loop automation. If a field technician needs to repair a connectivity problem, the controller can assist using computer vision and augmented reality.

Anomaly detection – the **sense** part

The controller continuously analyzes the data collected from the network. This analysis allows operators to recognize patterns, predict trends, correlate information, detect anomalies and prioritize network events. In this section, we will focus on how AI/ML can help automate and improve anomaly detection in a network. A seemingly simple approach to anomaly detection is to define a region of normal behavior and classify any observation outside this region as an anomaly. However, defining such a region is challenging. Moreover, this region typically evolves over time. As a result, numerous anomaly detection techniques exist.

Anomaly detection techniques can be broadly categorized into three classes: rule-based methods; statistical or model-based methods; and machine-learning-based methods. ML-based methods can be further divided into (semi-)supervised classification-based methods and unsupervised methods. In most scenarios, data is unlabeled, necessitating the use of unsupervised techniques, which do not rely on training data. Unsupervised methods are further subdivided into distance-based methods and clustering-based methods. More information about anomaly detection can be found in the surveys [7] and [8]. In the remainder of the section, we focus on the application of anomaly detection in fixed networks.

Rule-based thresholds

Rule-based (or heuristics-based) methods rely on predefined rules or thresholds to classify anomalies. In the context of fixed networks, a threshold crossing alarm (TCA) that is manually configured is a rule-based anomaly detection method. This type of TCA has drawbacks: it requires domain knowledge, manual effort, remains static without adaptation over time, and uses per indicator (e.g., board temperature) a network-wide generic threshold value that isn't adapted to the monitored object (e.g., a specific board). Model-based techniques offer solutions to overcome these drawbacks.

Methods based on prediction models

If the data consists of a time series, then model-based techniques that fit a prediction model to the time series are widely used. For streaming time series, such prediction-model-based techniques can assess whether a data instance is an anomaly immediately upon its arrival.

The simplest model is a fingerprinting model. This method automatically learns a threshold once at startup. One drawback of fingerprinting-based TCA is that the threshold remains static, failing to adjust to normal variations over time. To overcome this limitation, the model can be retrained using a moving-average (MA) approach. While an MA-based TCA adapts the threshold over time, its ability to predict normal variation is limited.

Network metrics, like those influenced by temperature or traffic, often show daily and seasonal patterns, complicating the detection of real anomalies. In this case, we need an ML-based prediction model. Typically, this involves fitting a regression model to the time series data. We propose using a decomposable time series model that includes components for trend, seasonality, and holidays. This model can detect anomalies in real-time as new data comes in. Furthermore, since the model forecasts the future trend, we can use a rule-based TCA to detect an anomaly well in advance. Time-series forecasting can be used not only for such preventive maintenance but also for other purposes, such as capacity planning. Figure 2 illustrates rule-based and model-based anomaly detection.

Figure 2. Rule-based and model-based anomaly detection.

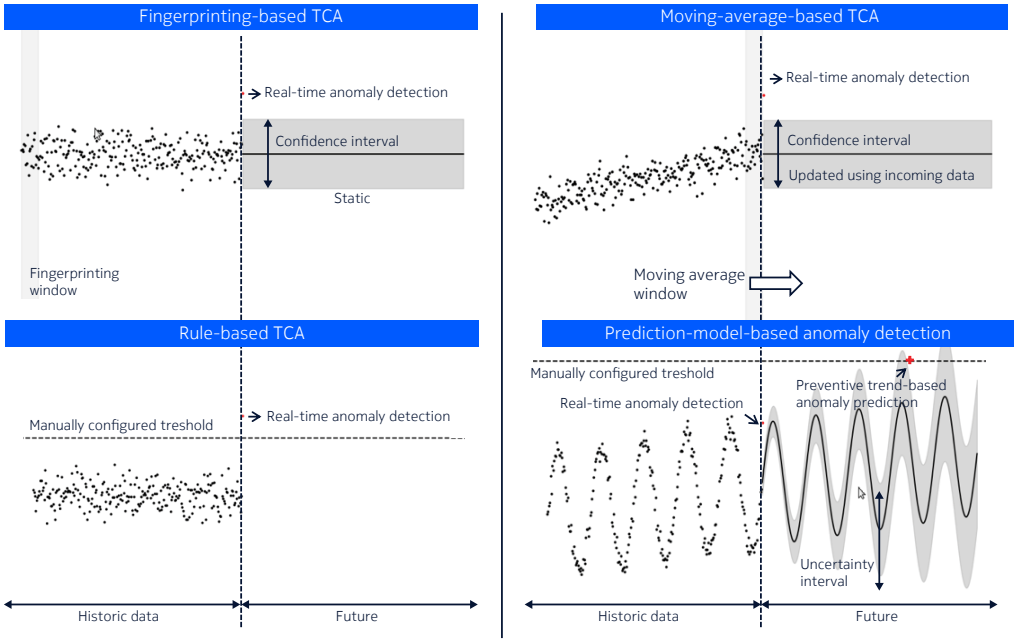

Model-based anomaly detection applies to diverse network metrics such as optical signal power, board temperature, CPU load, and RAM usage. An anomaly in the optical signal power may be caused by a fault in the outside distribution network (ODN) such as a fiber bend, a pressurized fiber, a loose, dirty, or improper connector, or a disconnected fiber. An increasing board temperature may be caused by a dusty filter.

Multivariate anomaly detection

An important component of a PON network is the SFP (small form-factor pluggable) transceiver at the OLT (optical line terminal) and ONT (optical network terminal). An SFP may degrade over time due to aging, connector wear, dust accumulation, high humidity, etc. To detect SFP degradation, multiple performance indicators are monitored: the optical power level, operating temperature, laser bias current, etc. However, SFPs of different vendors and types may have a different normal range for each of these metrics. Therefore, detecting a degraded SFP is challenging and requires machine learning.

With access to multivariate data, the preferred anomaly detection method is based on clustering. Machine learning groups similar data points into clusters, where each cluster represents a distinct combination of vendor and type of SFP. A data point that does not fit well into the cluster that corresponds to its vendor and type represents a degraded SFP. Such anomaly detection enables preventive maintenance.

We may also apply multivariate anomaly detection on a larger scale, at the level of an access node or OLT. The complexity of an access node involves monitoring so many metrics that detecting anomalies with traditional univariate rule-based methods is challenging.

Bandwidth usage anomalies

While on average a PON has abundant capacity to support tens of users during peak hour, the introduction of (multi-) gigabit services has led to PONs being at risk of congestion during noticeable time periods. Indeed, such high-speed services are used by subscribers to download or update video games up to several hundreds of gigabytes; or to download complete box-sets of television programs, for offline viewing. Therefore, long high-speed downloads are occurring more often, and users dominating the bandwidth for significant time periods are more common.

From a user's perspective, PON congestion risk is equivalent to the probability of a failed speed test. We can monitor speed test success probability continuously for every PON in the network using a digital twin network as we explain in the next section. Rule-based TCA is used to trigger a PON alarm when the speed test success probability drops below, for example, 80%. PON congestion may be caused by a single user dominating the bandwidth. Such an anomalous bandwidth usage can be detected using rule-based TCA as well.

Another type of bandwidth usage anomaly occurs when a user has limited bandwidth because of a bottleneck in the network. A bandwidth bottleneck can occur due to congestion at various locations in the network, for example, at the home Wi-Fi, at the PON downlink, and at the PON uplink. A bandwidth bottleneck can also be due to the speed limit of the service to which the user is subscribed. A bandwidth bottleneck anomaly can be detected by analyzing a user's bandwidth consumption over time using a statistical method.

Alarm correlation

While the controller can detect anomalies, so too can the OLT and ONT. Anomalies detected in OLTs and ONTs are reported back to the controller as alarms. The controller can correlate equipment alarms for root cause analysis and fault localization. Through alarm correlation, rogue ONTs can be detected and cut fibers localized.

Summary

Figure 3 maps every fixed-network anomaly that we described above to a network location.

Figure 3. Overview of potential anomalies in the fixed network.

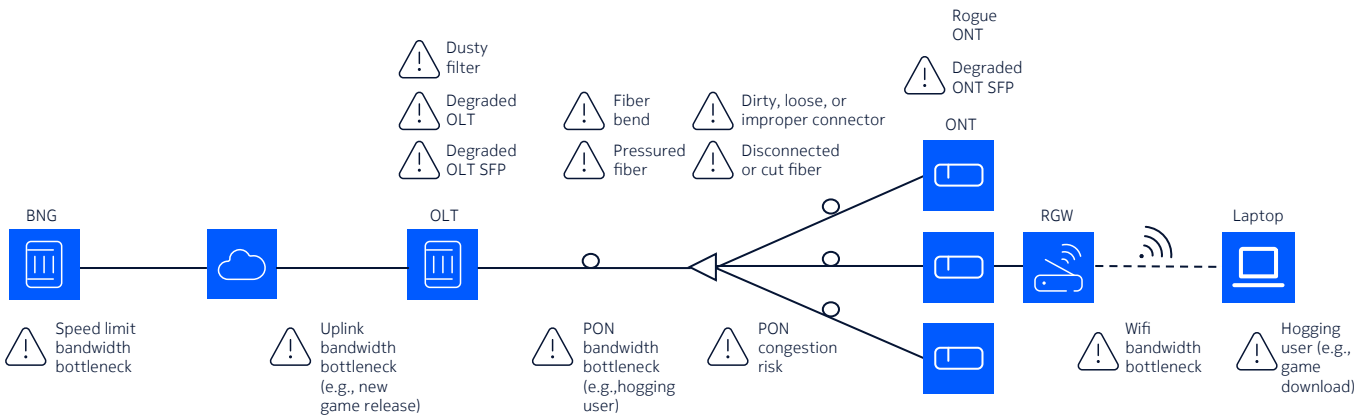

Digital twin network – the **think** part

The digital twin network can be utilized for analyzing, diagnosing, and controlling its physical network counterpart. It can improve planning decisions, validate configurations, speed up troubleshooting and help decide on the necessary corrective steps. In this section, we focus on the application of a digital twin network to improve the capacity management of a PON, both the strategic and operational aspects.

Strategic capacity planning involves defining service tiers and split ratios, considering future traffic growth. Operational capacity management includes finding the optimal PON to add a new user and verifying whether a user can be upgraded to a higher service tier without risking PON congestion. Moreover, DTN-based operational capacity management helps to resolve and even avoid the bandwidth usage anomalies described in the previous section.

Digital twin architecture

A digital twin network, a virtual representation of a physical network, is standardized in the ITU-T recommendation Y.3090. Figure 4 integrates the DTN architecture of Y.3090 (Figure 8-1 in the ITU-T recommendation) with the cognitive controller's sense-think-act framework (shown in Figure 1) for the specific application of integrated PON capacity management.

Figure 4. Digital twin network (DTN) integrated in a network domain controller.

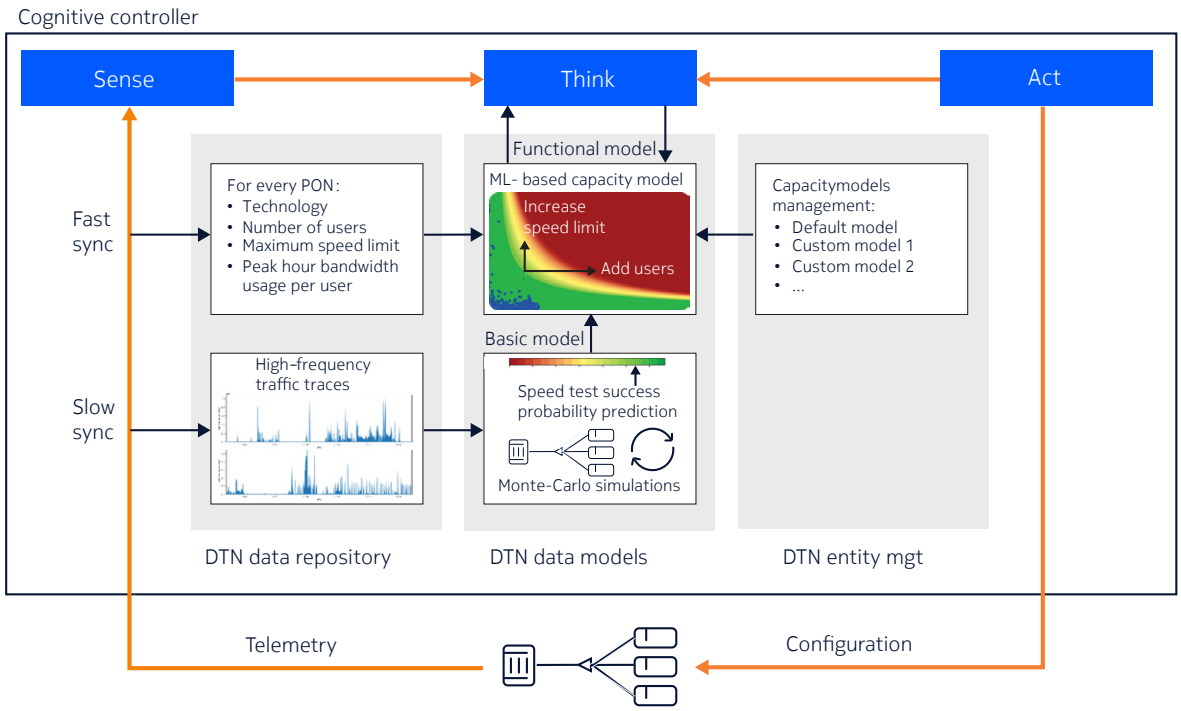

The digital twin network includes the three subsystems specified in Y.3090: the data repository, data models, and digital twin entity management. Because the DTN is integrated in the controller, for some of its functions required by Y.3090, the DTN relies on infrastructure provided by the controller such as the data collection, data repository, and interfaces. The data collection enables the real-time synchronization of the digital twin network with the physical network. This synchronization distinguishes a DTN from a traditional network simulator, enabling highly accurate data-driven capacity planning and operational capacity management.

DTN data models

The primary goal of PON capacity management is to ensure high availability of advertised speeds. Therefore, the DTN predicts speed-test success probability considering the capacity of the PON technology, the number of users connected, the highest service tier assigned to a user, the average peak hour bandwidth usage, and the presence of dominant users.

This prediction requires a representative collection of high-frequency traffic traces (with a 5-second sampling period), which capture the burstiness of internet traffic. Because the statistical properties of the internet traffic change only slowly over time as new services and applications are adopted, this collection of high-frequency traffic traces is needed only a few times per year (slow synchronization). The high-frequency traffic traces are used by the so-called "basic" model (Y.3090 terminology). This model simulates the PON traffic management (including traffic shaping and weighted fair queueing) to predict the speed-test success probability for many randomly selected PON and service configurations.

The large collection of simulation results enables the training of a physics-informed machine learning model, which serves as the Y.3090 "functional" model. This ML model provides robust and fast generalization of the simulated data and covers the complete capacity planning space as shown in Figure 5. The functional model enables the prediction of speed-test success probability for every PON (represented by blue) based on the fast synchronization of the DTN with the physical network. The speed-test success probability ranges from high (above 80% = green) to low (below 80% = red).

DTN use cases and benefits

Figure 5. Screenshot of the PON Capacity Planner Altiplano application, which implements a digital twin network for integrated PON capacity management. The configuration pane on the left enables the what-if analysis. The chart on the right shows the predicted speed-test success probability.

For strategic capacity planning, the DTN enables what-if analysis as illustrated in Figure 5. What is the impact if we increase the split ratio or if a (multi-) gigabit service is introduced. The DTN can also assess the impact of future traffic growth and of new PON technologies. DTN-based data-driven capacity planning is more accurate and QoE-focused than traditional capacity planning, avoiding overly conservative service definitions and PON configurations. In other words, the DTN enables an increase in PON utilization without risking more failed speed tests.

For operational capacity management, the DTN monitors the speed-test success probability for all PONs in the network continuously (blue dots in Figure 5) without the need for disruptive real speed tests. In IBN terminology, we state the intent to keep each PON's speed-test success probability above 80%. The cognitive controller verifies daily for each PON whether this intent is met. When a congestion risk alarm is raised, the controller may be able, in many cases, to autonomously bring the speed-test success probability back above 80% by smart network reconfiguration as we explain in the next section. For cases where an automatic resolution is not feasible, the DTN may suggest a user re-allocation or a PON technology upgrade. Prediction-model based anomaly detection may even identify PON congestion risks in advance.

For underutilized PONs, without a congestion risk, the DTN verifies whether adding a user to the PON or upgrading a user to a higher service tier is feasible without jeopardizing the required speed-test success probability as illustrated in Figure 5 (bottom table). Thus, the DTN enables intent validation. To minimize the impact on speed-test success probability, the DTN may even determine which PON to add a user to. The DTN may also predict the highest speed that can be offered to a user without risking PON congestion. In other words, the DTN also facilitates intent suggestion.

Closed-loop automation – the **act** part

While anomaly detection smartens the "sensing" and the digital twin network supports the "thinking", closed-loop automation enables the cognitive controller to take actions autonomously, without manual intervention. Closed-loop automation is directly related to closed-loop feedback control, a concept fundamental to control engineering, that also underpins intent-based networking as shown in Figure 6.

Figure 6. Intent life cycle (RFC 9315, Figure 1) rearranged to highlight that IBN enables closed-loop control.

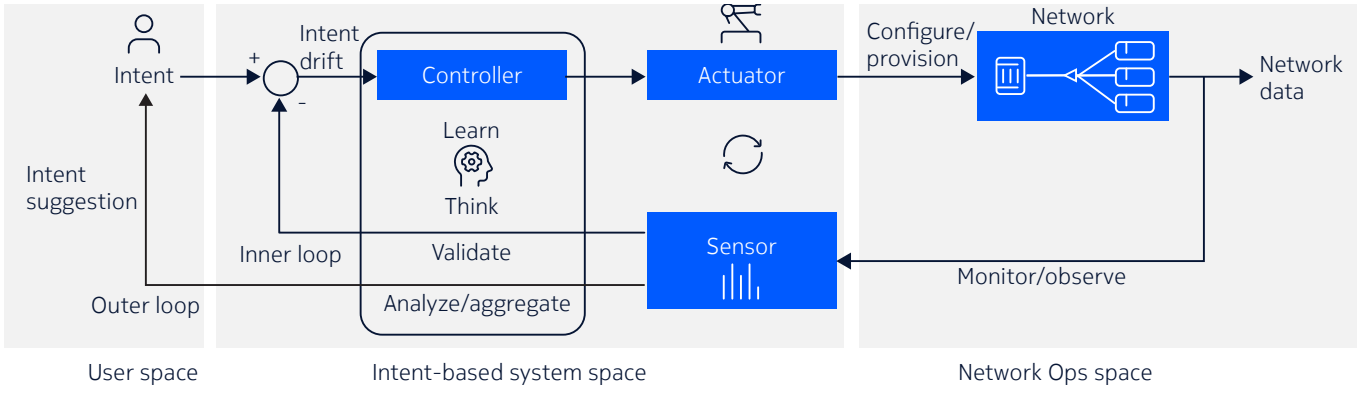

A closed-loop control system measures the actual output and feeds this signal back to compare it with the desired output. Intent-based networking implements such a closed-loop control as shown in Figure 6. The inner control loop corresponds to traditional closed-loop control. This fast-paced feedback loop ensures that the network behaves according to the intent expressed by the operator. The outer control loop feeds back to the user space, enabling intent optimization. While the inner control loop is fully automated, the outer control loop provides the (human) operator the option to select, postpone, ignore, or auto-execute the suggestions generated by data-driven analysis. The feedback may be relayed via a service orchestrator.

In the following sections, we focus on two use cases of closed-loop control. The first use case, which focuses on inner-loop control, automatically mitigates the congestion risk caused by a user dominating the PON bandwidth through dynamic traffic management. The second use case, demonstrating outer-loop control, automatically pursues service upgrade opportunities for users who experience bandwidth constraints due to the speed limit of their subscription.

Automatically mitigating congestion risks from hogging users

To ensure fair bandwidth sharing between subscribers in the downstream direction, the OLT implements weighted fair queuing (WFQ). Upon congestion, the PON's capacity gets fairly distributed between the active subscribers. Because a WFQ scheduler operates over short time periods and is memoryless, it does not prevent users from dominating the bandwidth over longer time periods. Therefore, CSPs often enforce a fair-use policy using a data volume cap. We propose an alternative method that mitigates congestion risks without a noticeable impact for the hogging user. This method is based on dynamic traffic management.

With dynamic traffic management, the sensor (Figure 6) measures periodically (typically every 5 minutes) the volume of data downloaded during a predefined time interval. The measured data volume is compared with a rule-based data-volume threshold (intent in Figure 6). If a hogging user is detected, then the controller determines the reduction of this user's weighting required to mitigate the congestion risk. The actuator applies this weighting reduction in the WFQ scheduler. The effect of such a preventive weighting reduction is illustrated in Figure 7.

Figure 7. Dynamic traffic management ensures speed test success for any user

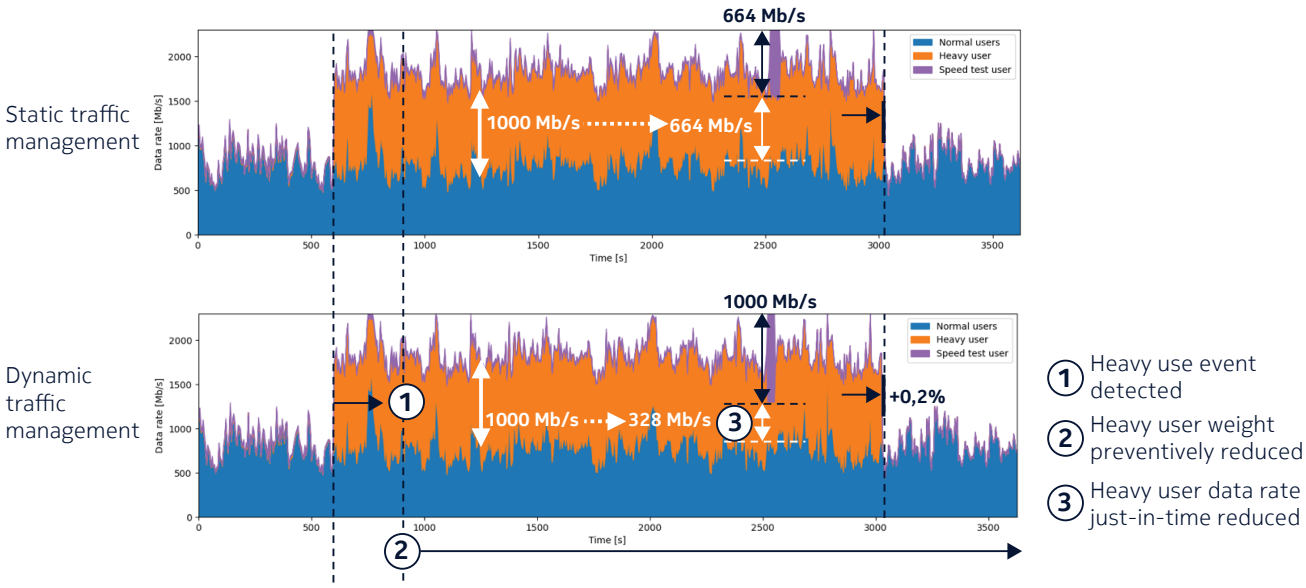

This figure compares static and dynamic traffic management for a PON with normal users (blue), a hogging user (orange), and a speed test user (purple). All users have a speed limit of 1 Gb/s. With static traffic management, which relies only on the WFQ scheduler, the speed test fails. Dynamic traffic management detects the hogging user (1), and therefore preemptively reduces its weighting (2). This weighting reduction does not affect the hogging user's speed except during the short period of the speed test when the PON is congested (3). With dynamic traffic management, the speed test is successful while the increase in download duration for the hogging user is negligible.

Automated pursuit of upgrade opportunities for constrained users

Users who experience bandwidth constraints due to the speed limit of their subscription are likely to be interested in an upgrade to a higher-speed service. As explained in the section on anomaly detection, a bandwidth bottleneck can be detected by analyzing a subscriber's bandwidth usage over time.

If the bottleneck is due to the speed limit of the user's subscription, the controller can determine the speed required to mitigate the bandwidth limitation and verify that the required rate increase does not risk PON congestion. This verification is done using the digital twin network.

If the speed increase is feasible, the controller can provide an intent suggestion to the operator (OSS/BSS) using the outer control loop shown in Figure 6. The intent suggestion may be used to offer the subscriber a free trial of a higher-speed service. If the user accepts the offer, the operator adapts the intent. Or the operator may revert the intent if the user decides not to subscribe to the higher-speed service.

Assistance to humans

As network automation continues to advance, there remains a crucial role for human involvement in specific aspects of operations. Even with automated processes handling routine tasks, humans still need to ingest intent and troubleshoot problems that the controller cannot autonomously resolve or diagnose. In addition, human technicians are still necessary for conducting manual interventions in the field, particularly for installing and maintaining the outside distribution network. AI/ML paves the way to assist with human interventions as we explain in the following sections.

Computer vision

Computer vision is a branch of artificial intelligence focused on enabling machines to analyze, process, interpret, and understand visual data from images and video. While computer vision began developing in the late 1950s, it has significantly advanced since 2012 due to the convergence of three key factors: the introduction of deep learning models, the unprecedented availability of labeled images for model training, and the dramatic increase in the computational power of GPUs. Today, accuracy of image recognition rivals human performance and even surpasses it in specialized applications.

Computer vision can assist field technicians in deploying and maintaining the outside distribution network (outer control loop 2 in Figure 1). Each technician is equipped with a smartphone featuring a high-quality camera, enabling them to capture images of the network elements before and after manual operations. The images are sent for analysis using object detection, image recognition, optical character recognition (OCR), barcode recognition, and data matrix code recognition. The goal of this OCR and code recognition is to uniquely identify any labeled network element handled by the technician to ensure that the technician works on or with the right element and to keep the inventory management system error free.

The object detection and image recognition enable instant detection of faults in the physical network installation and configuration. Because faults are detected while the technician is on site, we avoid unnecessary truck-rolls and network deployment delays. Augmented reality is used to display both the identified faults and suggested repair instructions. Technicians can see visual overlays on the image they captured that highlight the faults and guide them through the necessary repairs. By automating the quality control process, field operations can be monitored and verified in real-time. This helps improve the accuracy of the installation from the start, reducing errors and increasing the rate of first-time right operations.

For example, when a technician needs to connect a new subscriber, the controller identifies the recommended splitter port and, if possible, one or more alternative splitter ports. This recommendation is based on information from the inventory management system and passive topology database. Before making the connection, the technician captures a first image to verify the current state of the drop splitter installation. The controller analyzes this image and instantly provides feedback to the technician using augmented reality. In the example shown in Figure 8, the controller flags that the recommended port is already in use and three of the ports lack termination caps. After the installation, a second image is captured and sent to the controller for verification. If the installation is faultless, then the inventory management system and passive topology database can be updated.

Figure 8. Example of how AI helps a field technician connect a new subscriber to an existing drop splitter.

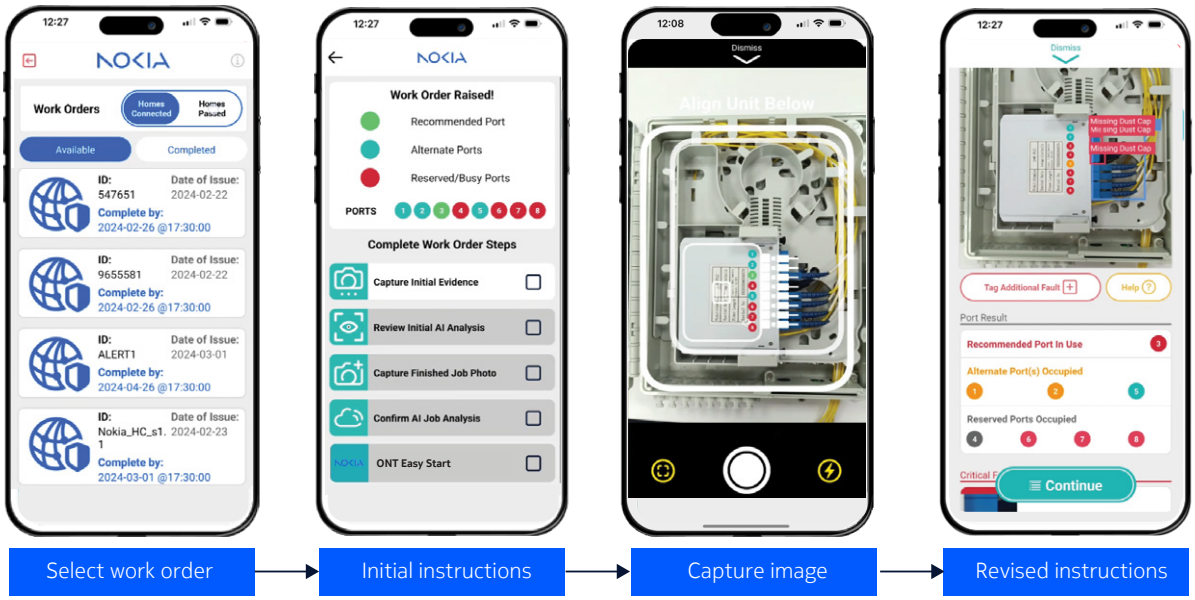

Computer vision supports capacity management of the outside plant (ports in use, ports free in street cabinets, distribution points, splitters, etc.), helps identify the optical path of a subscriber (to which port, splitter, patch cord the end user is connected), and proactively signals frequently observed installation errors (missing screws, dust caps, bent fiber, wrong patch cord, etc.).

Large-language models

The success of deep learning for computer vision inspired similar breakthroughs in natural language processing (NLP), culminating in the development of the transformer architecture. This innovation greatly enhanced the accuracy of large language models, making their capabilities comparable to human performance in many tasks. As a result, the public release of ChatGPT in 2022 showcased these models' ability to understand and generate human-like text with impressive proficiency.

LLMs can support operators in troubleshooting network problems. While a pre-trained LLM base model contains a wealth of information, such a model is not effective for such specialized tasks without additional techniques to augment its capabilities. To effectively troubleshoot network problems, the LLM needs to be fed with information from sources such as equipment documentation, log files, alarm notifications, and telemetry data. To provide the LLM with data relevant to the user query, the LLM is combined with an information retrieval system. This combination enables retrieval-augmented generation (RAG). The information retrieval system may use an LLM as well.

We are conducting research on such RAG systems. We have developed an experimental RAG system to allow operators to ask any question about the ONTs connected to their network. In a first step, the information retrieval system translates the operator's question, expressed in natural language, into a database query and then searches the ONT database to retrieve the relevant data. In the next step, we instruct the LLM to answer the question using the retrieved data. Figure 9 shows a screenshot of the user interface of this RAG-based chatbot. In addition to providing the answer, this chatbot also displays the database query used to obtain the relevant information. This feature enables the user to verify the accuracy of the query. The benefit of this chatbot is that users can ask questions without having to know database query languages or the database schema. Another advantage is that the system can handle a wide variety of questions without requiring predefined queries for each type of question.

Figure 9. Screenshot of an experimental LLM-based chatbot.

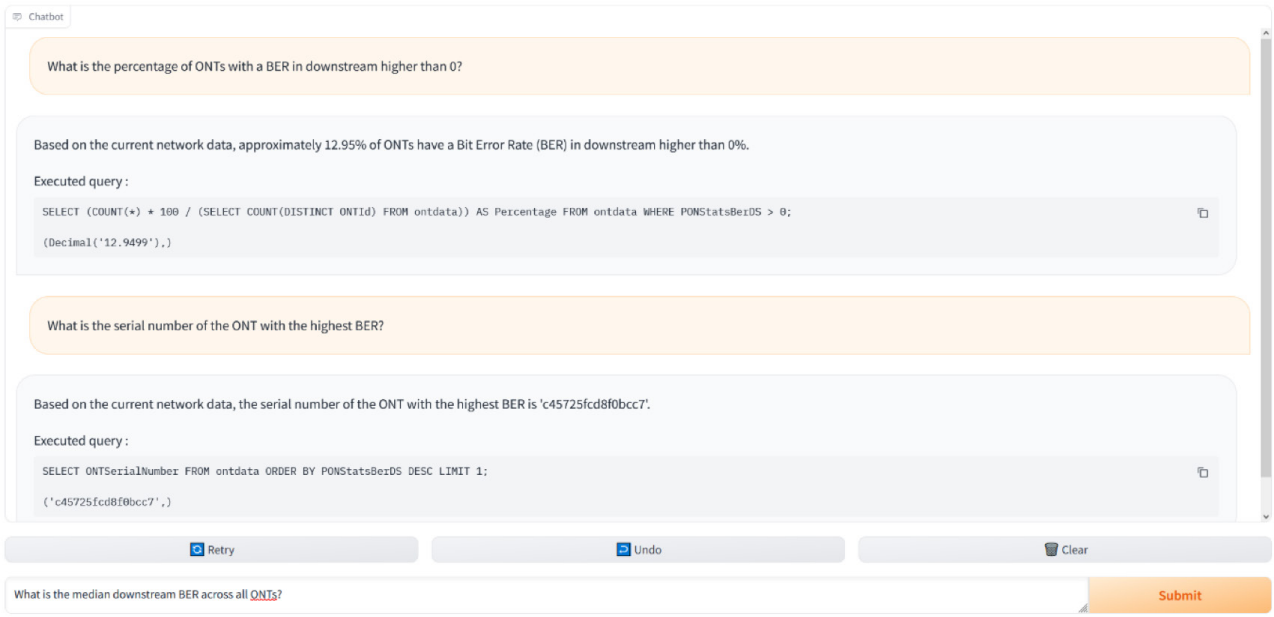

Another forward-looking application of LLMs is intent ingestion (outer control loop 2 in Figure 1). Intent ingestion is the process of obtaining intent statements through interactions with users. By using an LLM-powered chatbot for the intent ingestion, network operators can specify their intent using natural language. A human-machine dialog is used for intent clarification and refinement. The LLM allows operators to express intent without having to learn the specific language of the intent-based system, thereby avoiding a steep learning curve and reducing the required skill level. As part of this process, the LLM can use the digital twin network to either validate any intent expressed by the operator or make alternative intent suggestions.

Conclusion

Modern network domain controllers facilitate AI/ML-based decision-making, paving the way for autonomous broadband networks that can sense, think, and act. In this paper, we've presented numerous beneficial applications of AI/ML in fixed access networks. Most of these applications are already available for deployment today (see Appendix). Moreover, AI/ML is advancing more quickly than most other recent technologies. Broadband operators are well-positioned to take full advantage of these developments.

Appendix: mapping applications of AI/ML in fixed networks to Nokia solutions

Nokia's solution for the fixed-network domain controller is the Altiplano access controller. This controller is an open and programmable platform that enables data analytics and AI/ML diagnostics. It can be flexibly extended with applications that help to better optimize, troubleshoot, and analyze the fixed network.

Table 2. Mapping AI/ML applications and use cases to Nokia solutions.

Use case or AI/ML application	Technology	Nokia solution
Manually configured TCA	Anomaly detection	Altiplano feature
Fingerprinting-based TCA	Anomaly detection	Altiplano feature
Moving-average-based TCA	Anomaly detection	Altiplano feature
Real-time anomaly detection with time-series forecasting	Anomaly detection	Network Trend Analyzer (NTA)
Trend-based anomaly detection with time-series forecasting	Anomaly detection	Network Trend Analyzer (NTA)
Detection of degraded SFP	Anomaly detection	SFP Health Monitor (SHM) and ONT Health Monitor (OHM)
Trend-based congestion risk detection	Anomaly detection	PON Capacity Planner (PCP)
Detection of a hogging user	Anomaly detection	Bandwidth Sharing Optimizer (BSO)
Detection of bandwidth bottleneck	Anomaly detection	Service Campaign Manager (SCM)
Detection of rogue ONT	Alarm correlation	Automated Troubleshooting Assistant (ATA)
Localization of disconnected or cut fiber	Alarm correlation	Automated Troubleshooting Assistant (ATA)
Strategic capacity planning • What-if analysis • Split ratio optimization • Service tier definition	Digital twin	PON Capacity Planner (PCP)

Use case or AI/ML application	Technology	Nokia solution
Automatically mitigating congestion risk from dominant users	Closed-loop automation	Bandwidth Sharing Optimizer (BSO)
Automated pursuit of upgrade opportunities for constrained users	Closed-loop automation	Service Campaign Manager (SCM)
ODN installation assistance • Installing a new splitter • Connecting a home to an existing splitter • Map splitter ports to subscribers to refine passive topology	Computer vision	Broadband Easy Connect

References

1 Automation and AI/ML in fixed networks, Nokia White Paper, 2021

2 Software-defined access networks, Nokia white paper, 2023

3 Broadband network telemetry, Nokia White Paper, 2023

4 Intent-Based Networking - Concepts and Definitions, IRTF RFC 9315, 2022

5 Autonomous Networks Technical Architecture, TM Forum IG1230, 2023

6 Digital twin network – Requirements and architecture, ITU-T Y.3090 recommendation, 2022

7 Zero-touch network and Service Management (ZSM); Network Digital Twin, ETSI GR ZSM 015, 2024

8 Anomaly Detection: A Survey, V. Chandola et al., ACM Computing Surveys, 2009

9 A Review on Outlier/Anomaly Detection in Time Series Data, A. Blázquez-García et al., ACM Computing Surveys, 2021

10 Zero-touch network and Service Management (ZSM); Closed-Loop Automation, ETSI GR ZSM 009, 2023

11 Detect and predict broadband issues with Altiplano Network Trend Analyzer, Nokia Blog, 2023

12 How data is driving the future of PON capacity planning, Nokia blog, 2023

13 Assure peak bandwidth for broadband subscribers with new Altiplano app, Nokia blog, 20224

About Nokia

At Nokia, we create technology that helps the world act together.

As a B2B technology innovation leader, we are pioneering networks that sense, think and act by leveraging our work across mobile, fixed and cloud networks. In addition, we create value with intellectual property and long-term research, led by the award-winning Nokia Bell Labs.

Service providers, enterprises and partners worldwide trust Nokia to deliver secure, reliable and sustainable networks today – and work with us to create the digital services and applications of the future.

Nokia OYJ
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

(September) CID214180