Independent market research and competitive analysis of next-generation business and technology solutions for service providers and vendors

HEAVY READING
WHITE PAPER

AI/ML for CSP Operations: From Reactive to Predictive & Autonomous

A Heavy Reading white paper produced for TEOCO

TEOCO

AUTHOR: JAMES CRAWSHAW, SENIOR ANALYST, HEAVY READING

INTRODUCTION

The complexity of communications networks seems to rise inexorably with the deployment of new services, such as software-defined wide-area networking (SD-WAN); new network architectures, such as 5G; and new technology paradigms, such as network functions virtualization (NFV). To meet ever-increasing customer expectations, communications service providers (CSPs) need to increase the intelligence of their network planning and operations. To do so, CSPs are investing in new analytics capabilities that take advantage of the mathematical domains of artificial intelligence (AI) and machine learning (ML).

The terms AI and ML are often used interchangeably. In Building the Network of the Future, Mazin Gilbert and Mark Austin's chapter ("Network Data and Optimization") defines ML as "a computer program [or software agent] that can learn from data to predict a future state or condition." In contrast, Gilbert and Austin define AI as "a more advanced form of an ML system that is able to act and continuously optimize its actions over time and context." While ML is like predicting the best next move in chess, AI is like formulating a strategy to win the game and then updating this strategy with every turn.

Gilbert and Austin categorize analytics into four camps, as shown in **Figure 1**. These different categories are an evolution from describing the network state, to reacting to failures, to predicting problems, and finally to taking automated actions to avoid predicted problems. We believe CSPs will undertake a similar adoption curve with AI/ML in network management. In the initial phase (descriptive and reactive), they will use AI/ML systems to observe events and make correlations that were previously outside their view. Gradually, they will rely on AI/ML systems to predict incidents and make recommendations on how to resolve them. And finally, the systems will be allowed to operate with increasing autonomy, applying their recommended changes without the need for human intervention.

Figure 1: Analytics Classifications – From Passive to Proactive

Type	Description
Descriptive Analytics	Summarizing data and presenting visually in a table or chart to monitor a system or service.
Reactive Analytics	Raising alarms based on events such as hardware failures. Uses traditional rules and statistical methods such as clustering.
Predictive Analytics	Predicting events and future states based on time series data. Can use supervised, unsupervised or semi-supervised learning. Algorithms can be linear (e.g., regression) or non-linear (neural networks, support vector machines).
Prescriptive Analytics	Similar to AI, where the goal is to identify an optimized sequence of steps or actions to take. An example would be optimizing data routing across all network nodes and end devices as traffic changes.

Source: Heavy Reading, based on Gilbert and Austin's "Network Data and Optimization"

APPLYING AI/ML TO THE NETWORKING DOMAIN

AI/ML is a key enabler for automation. Many telecom operators have already implemented AI/ML solutions as part of their customer care and marketing armory. Applying AI/ML to the networking and service domains has proved to be tougher, although approaches are now beginning to emerge. Some of the key drivers behind the adoption of AI/ML-based analytics in networking include:

- **Cost efficiency:** The need to improve performance and QoE without increasing cost, as subscribers are more demanding but less willing to increase their spending
- **Usage growth:** Traffic growth outpacing traffic capacity, and becoming more differentiated in terms of applications, services and devices
- **New technology:** Network evolution (from 4G to 5G, from appliance-based to virtualized and software-defined) benefits from a more extensive use of analytics

Generic AI/ML frameworks can be useful for developing applications across a wide range of business activities, from the call center to the retail outlet. But for specialist areas such as telecom networks, there is a need for more focused AI tools that use specialized algorithms and have been trained with the relevant data sets. We can't just feed data into an algorithm and hope that useful insights pop out. Instead, we must apply a rigorous, scientific method to AI/ML, which starts with clearly defined problems, moves on to testable hypotheses, and then designs experiments to test those hypotheses.

For network analytics, we need to have telecom expertise embedded in the solution, as this knowledge is vital to the generation of the right mathematical models and to the constant tuning of their parameters. For example, when generating a predictive ML algorithm to predict network events, while the algorithm can be executed to predict various alarms or future failures, it only makes sense to focus on the predictions that will actually trigger a proactive action and can really help improve the network operation process. This kind of domain expertise would typically come from a network operations specialist, not an algorithm developer.

Part of the challenge is that, unlike image and voice recognition, there is no "Theory of Networking" around which to build an AI optimization algorithm. As Brian Levy points out in How Will AI and Machine Learning Impact CSPs?: "One of the biggest challenges when applying machine learning for network operation and control is that networks are inherently distributed systems, where each node (for example, switch, router, etc.) has only a partial view and partial control over the complete system. Learning from nodes that can only view and act over a small portion of the system is difficult and complex, particularly if the end goal is to exercise control beyond the local domain."

APPLICATION OF AI/ML TO NETWORK OPERATIONS

AI/ML could support network operations in a number of ways. Initially it might be used to help detect issues – such as faults or service-level agreement (SLA) breaches – in real time or near-real time. A more advanced use case would be for diagnosing root causes of alarms, correlating across multiple event sources and filtering out noise (false alarms). Finally, we might employ AI/ML algorithms to recommend solutions to the root cause we have identified. Although some of these capabilities are built into existing service assurance solutions, they may struggle with the move to 5G and associated technologies such as NFV, due to increased levels of abstraction in the network design, which complicate correlation analysis.

ML applications for network operations will make use of existing service assurance data sources such as alarms, device measurements and customer usage records. Using multiple data sources can help to make better predictions, as each data source will reflect a different aspect of network and service behavior. The data is then analyzed using various ML techniques such as regression modelling, clustering, Bayesian graph models, and neural networks, as outlined in **Figure 2**.

Figure 2: Machine Learning Mindmap

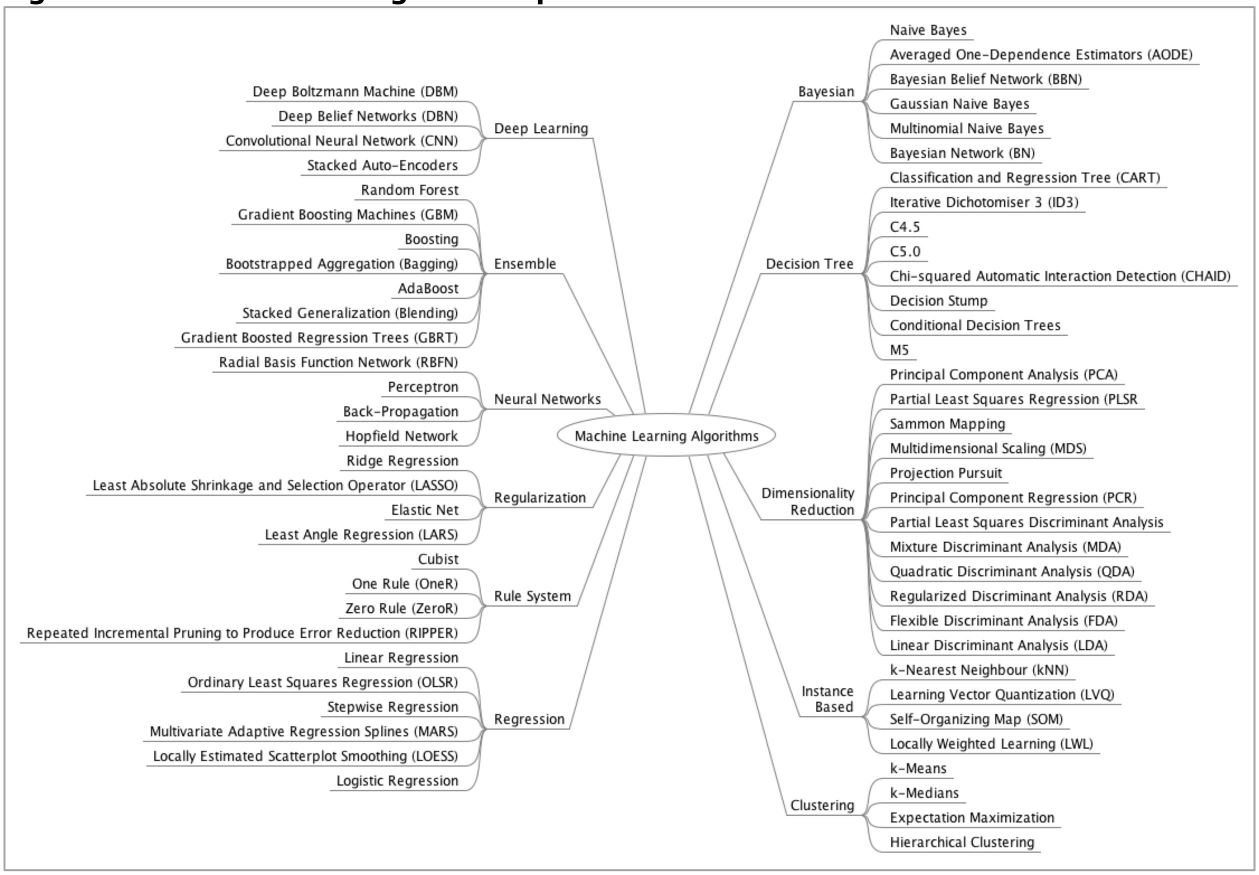

Source: Machine Learning Mastery

Below we consider some potential areas for the application of AI/ML in network and service operations.

Anomaly Detection

With AI/ML algorithms, operators can look for correlations between KPIs and faults, and also between the KPIs themselves in order to identify leading indicators. While humans might be able to track a handful of KPIs and their interrelationships on a dashboard, AI/ML can enable thousands of such relationships to be tracked in real time, while discovering "non-trivial" KPI relationships and abnormal behavior.

Another opportunity for AI/ML is automatic adaptation of KPI thresholds. Typically, KPI thresholds are set once and when breached, an alarm is generated. However, different statistical algorithms could be used to calibrate and fine-tune the thresholds to detect non-trivial problems at an early stage and adjust to dynamic changes in network behavior. Automated threshold creation and adjustment would reduce the time to deploy new services and technologies, the time to detect and recover from network degradations and eliminate false positives typically generated by manually created and maintained thresholds.

Alarm Prioritization & Root-Cause Analysis

Network operations staff are bombarded with so many alarms, they struggle to see the forest for the trees. AI/ML can help by classifying alarms according to historical behavior, and by identifying their root cause. Operators can then suppress the symptomatic alarms and focus on the root cause alone, thereby eliminating redundant trouble tickets and unnecessary remediation efforts, along with their associated costs.

AI/ML can extend traditional rule-based root-cause analysis (RCA) with adaptive mechanisms to locate the source of problems more quickly. AI/ML algorithms analyze the alarms, automatically suggesting groupings and correlations and tagging the potential root-cause alarms among them, or creating synthetic alarms that describe the cause of the failure.

AI/ML could use clustering to find correlations between alarms that had previously been undetected or use classification to train the system to prioritize alarms. Traditional rule-based alarm correlation suffers from a heavy burden of rule maintenance. With AI/ML, we could instead train a system to devise its own rules, based on a given set of data inputs. AI/ML-based RCA could automatically derive the relationship between network elements and events without the need for predefined rules.

Figure 3 shows the real-world results of AI/ML-based alarm prioritization at two different operators. Based on automated analysis of alarm history and the related user behavior and actions, the system assigned each of the current active alarms one of three tags to mark its importance: Important, Standard or Ignore.

Figure 3: Automated AI/ML-Driven Alarm Prioritization - Real-Life Results

Operator	Tier 1 Wireline	Tier 1 Mobile
Active alarms	35,000	185,000
Important	8%	20%
Standard	74%	59%
Ignore	18%	21%

Source: TEOCO

These tags help network operations staff focus their attention on the most important alarms and shorten the time it takes to resolve the more critical issues. As can be seen, both operators were dealing with a very large number of active alarms at any point in time. In both cases, around 20 percent of the alarms could be suppressed thanks to the AI/ML-based solution. The system was able to help operators focus their efforts on 10-20 percent of the total alarms, which it deemed important. Applying the Pareto principle, these 10-20 percent of alarms were probably responsible for 70-80 percent of the network issues.

Action Recommendations & Self-Healing

ML could be applied to automate the resolution of common incidents that generate trouble tickets. The system could be taught by operations staff how to handle these incidents but still require human approval before taking action (supervised or open-loop operation mode). Longer term, as humans become more comfortable with AI/ML-enabled decision-making, they may let it operate with increasing autonomy. The phased adoption of AI/ML might look like **Figure 4**.

Figure 4: Phased Adoption of AI/ML in Telecom Networking

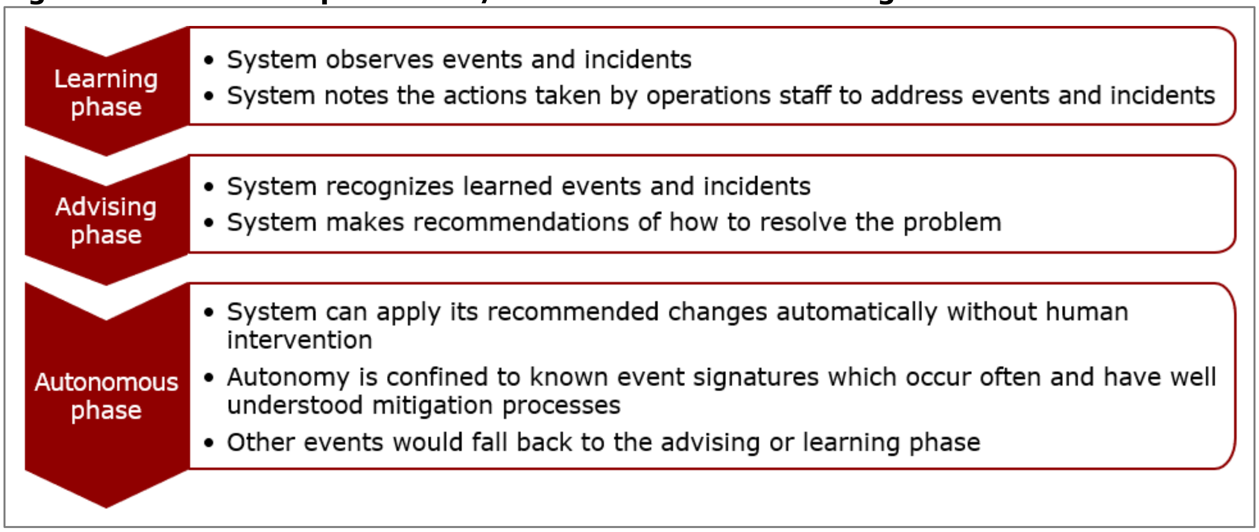

Source: Heavy Reading

Prediction of Network Faults & Predictive Maintenance

AI/ML algorithms can be used to make fault predictions by looking at trends (e.g., change in the volume of alarms over time), anomalies (versus peer-group entities) and causality analysis (e.g., event A causing event B). In a survey of CSPs, we found that predictive maintenance was the top use case for AI/ML in telecom, ahead of security, network management and fraud/revenue assurance, as shown in **Figure 5**.

Figure 5: Use Cases for AI/ML in Telecom

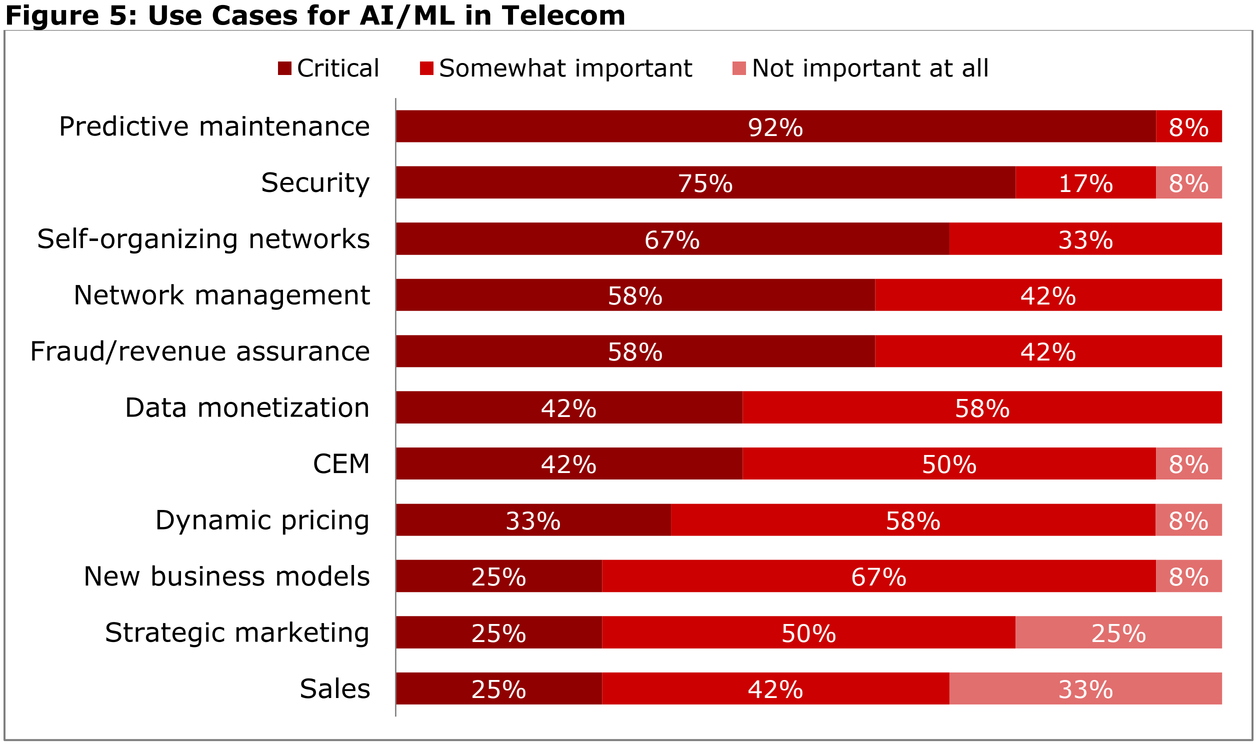

Source: Heavy Reading Survey – Thought Leadership Council (n=12), Nov. 2017

Anecdotally, we have found that senior executives at CSPs are excited about the prospect of fault prediction using AI/ML. However, network operations center managers will still face tough decisions, even with the help of these new tools. For example, if the system predicts that a fault will occur within the next month with 80 percent probability, does the manager replace the associated element immediately, make sure a spare unit is available in stock or wait until a failure occurs and scramble to fix it? While the first option ensures the greatest reliability, it is also the most capex-intensive. As the available technology matures, the corresponding network operation processes and culture need to evolve as well.

Figure 6 shows the results of a preliminary experiment conducted by AT&T using ML to predict network failures in a closed-loop system. Around 50 signals were used that provided KPIs of a network performance. The failure relates to a CPU overload and system shutdown. As can be seen, the closer to the failure event, the higher the accuracy of the prediction. But even going out as little as 50 minutes, the accuracy of prediction falls to 80 percent. The window of visibility in this example is fine for hot switchovers between servers, but for predictive maintenance of network infrastructure, a longer range of prediction would be required.

Figure 6: Predicting Network Failures in Advance

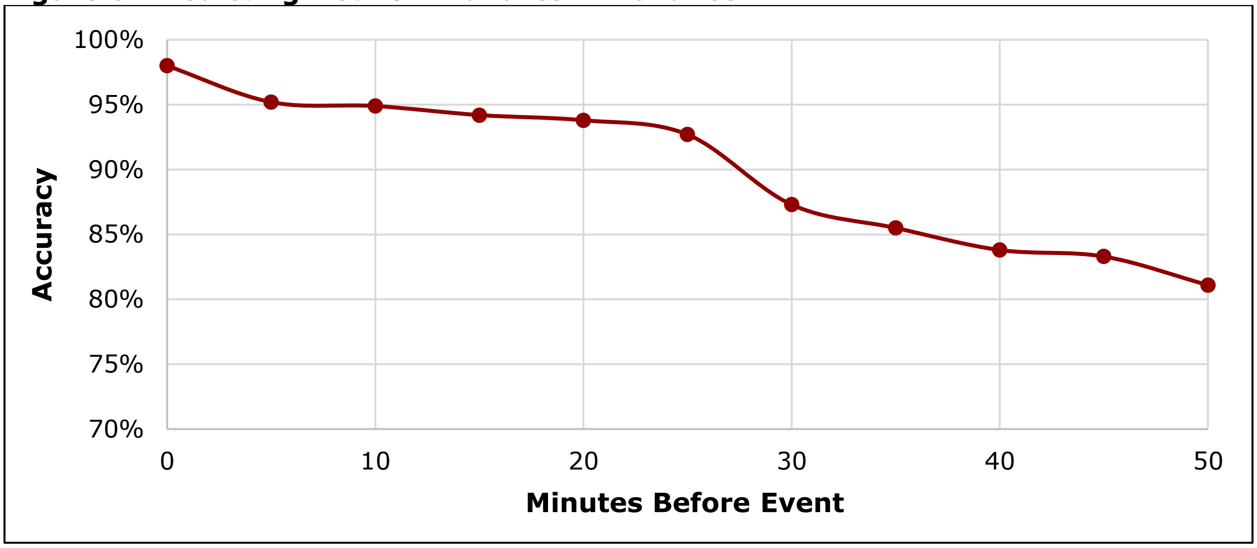

Source: Heavy Reading based on Gopalan's 2016 AT&T internal white paper, "Predictive ML for SDN Applications"

Predictive failure recognition could be extended beyond the network element level to making prediction of an end-to-end service availability at a generic level, or even at a specific customer level. This might help the operator avoid SLA penalties or customer churn by predicting service KPIs over time and notifying in case of potential service-level objective violations, so that preventive maintenance may be undertaken.

FINANCIAL JUSTIFICATION FOR AI/ML INVESTMENTS

The return on investment from applying AI/ML to network operations will generally come from a combination of the four factors outlined in **Figure 7**: lower opex, lower capex, higher service availability (avoiding SLA breaches and lost revenue opportunities) and faster repairs (which leads to better customer experience).

Figure 7: RoI From AI/ML in Networking

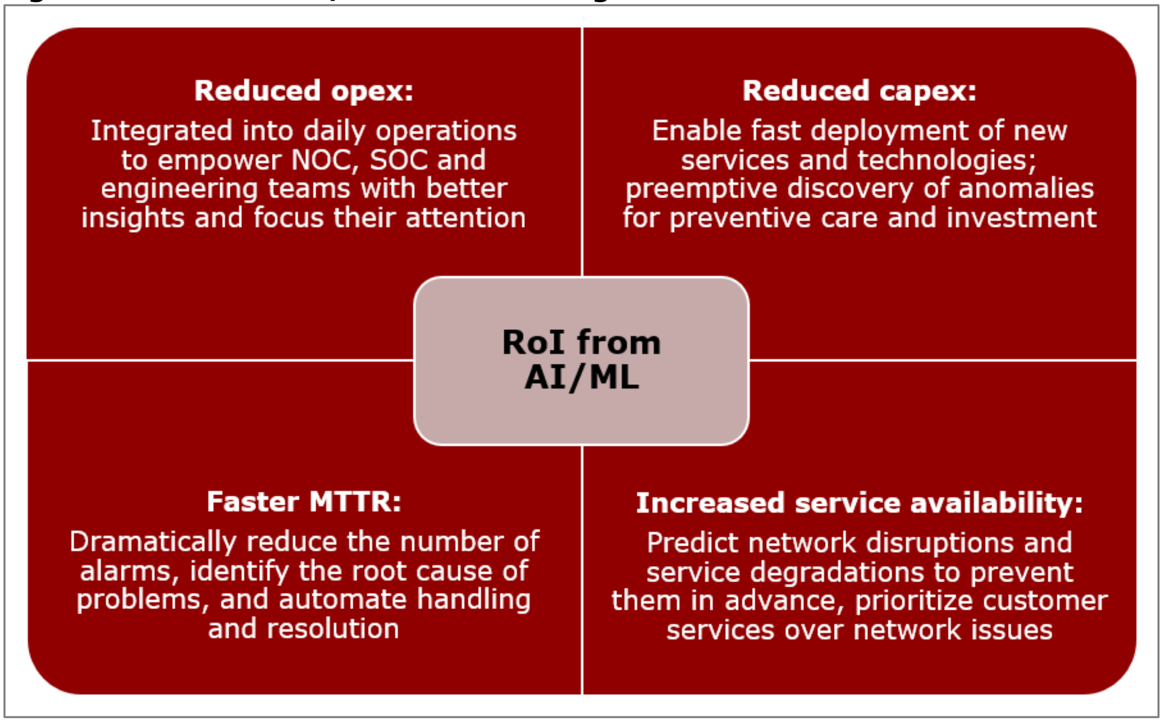

Source: Heavy Reading

CONCLUSIONS

Increased complexity in networking and networked applications is driving the need for greater network automation and agility. The advancement of NFV/SDN and transition to 5G will make monitoring and troubleshooting of the network a lot harder. The sheer volumes of data that network operations staff will need to handle require AI/ML-based solutions, as human-driven network and service management is becoming increasingly unsustainable. The emphasis is shifting from visualizing network performance and faults to feeding automated systems with insights and action points.

Network automation platforms such as the Open Networking Automation Platform (ONAP) should incorporate AI/ML techniques to deliver efficient, timely and reliable management operations. However, CSPs are not just looking for tools to manage these new technologies in isolation. Instead, they want tools that they can apply to a hybrid of legacy and new technology to manage the whole more reliably and with a greater degree of automation.

AI/ML brings new capabilities beyond traditional analytics, such as the ability to spot previously hidden patterns, detect unknown network anomalies and make predictions. With advances in computing and mathematics, these techniques are now affordable and practical enough to be applied in the data-rich telecom domain. In this paper we have discussed how AI/ML could be used for anomaly detection, alarm prioritization, root-cause analysis, remediation action recommendations, self-healing, prediction of network faults and predictive maintenance. Other applications include SLA validation prediction, network capacity planning (congestion prediction), and performance monitoring and optimization.

In general, AI/ML is suited to narrowly defined tasks, within a specified context and with consistent input variables. We are a long way from the sort of general intelligence that humans possess to learn from one situation and apply to another, addressing multiple problems and reacting to different inputs and changing scenarios. Network operations is a multi-faceted process that current ML technology can only partially address. The unpredictable nature of complex systems such as telecom networks requires the flexibility of general AI, which is a fascinating topic for academic research, but still some way away from being realized. Instead, we need to look at AI/ML as a tool to be embedded in the operations workflow to augment personnel, not just replace them.

Furthermore, AI/ML-based analytics is not a one-time fix for all our network ills. Instead, we should consider AI/ML as part of an ongoing process of continuous improvement in network management with the goals of improving the customer experience and keeping costs under control to ensure profitable growth for the business. In the initial phase, AI/ML might support existing processes by doing a better job of spotting anomalies and determining the root cause of alarms. As these systems learn over time, they may be able to help operations teams predict problems in the network ahead of time, so that they can undertake predictive rather than reactive maintenance. Finally, as our confidence in AI/ML systems grow, we will see them operate more autonomously, resolving problems in the network without human intervention, so that operations staff can focus on more intractable problems that AI/ML systems have not yet learned to master.

ABOUT TEOCO

TEOCO (www.teoco.com) is a leading provider of analytics, assurance and optimization solutions to over 300 CSPs and OEMs worldwide. Our solutions enable the digital transformation of CSPs while enhancing their network QoS, improving their customer experience and reducing their operational costs. Through advanced analytics and automation, TEOCO solutions provide actionable and measurable insights into network and customer behavior. This includes the optimization, effective monetization, and delivery of new and existing services, such as VoLTE and Video. Our commitment to network flexibility and agility makes TEOCO the obvious choice for CSPs looking to leverage NFV/SDN and the rise of 5G, and to maximize the revenue potential of new opportunities tied to video and the emerging Internet of Things (IoT).

Results-Driven Relationships

Industry-leading companies choose TEOCO because of our world-class solutions and unmatched commitment to measurable results. Our global customer base spans mobile, fixed, hybrid and next-generation networks.

Our success comes from our clients' success – we help to increase profitability and create more efficient networks. The powerful combination of technology and telecom-focused experience makes us both a market leader and a robust business management partner.

Extensive Product & Solutions Portfolio:

- **Service Assurance:** Delivers CSPs an automated, analytical, proactive and unified approach to service quality assurance, network performance management and fault management. By simplifying CSP operations and supporting the digital transformation, we enable the quick adoption of next-generation telecommunications services driven by technology advancements such as virtualization, NFV/SDN, 5G and IoT.
- **Business Analytics:** Provides a granular view of subscriber, network and business transactions, enabling CSPs to increase revenues, expand margins and maximize subscriber value and profitability. TEOCO touches all aspects of a CSP's business operations and provides insights into profitability, margins, roaming, costs, traffic, usage, subscribers, routes and partners.
- **RAN Solutions:** Helps engineers plan, optimize and manage a network to deliver maximum network performance and efficient capex utilization. TEOCO's RAN products and solutions are multi-technology and multi-vendor, delivering best-in-class planning, geolocation, audit and optimization capabilities to improve operational efficiencies and reduce total cost of ownership.
- **Device & IoT:** Provides device testing, monitoring and benchmarking solutions to the world's leading device manufacturers, CSPs and regulators to ensure mobile and IoT devices meet minimum performance requirements. For network operators, device certification is essential in ensuring subscriber satisfaction. Our state-of-the-art lab assists customers in the increasingly complex and time-consuming task of device certification.