HEAVY READING REPORTS

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

AI in Telecom Operations: Opportunities & Obstacles

AUTHOR: JAMES CRAWSHAW, SENIOR ANALYST, HEAVY READING

TABLE OF CONTENTS

1. EXECUTIVE SUMMARY

The complexity of communications networks seems to increase inexorably with the deployment of new services, such as software-defined wide-area networking (SD-WAN), and new technology paradigms, such as network functions virtualization (NFV). To meet ever-rising customer expectations, communications service providers (CSPs) need to increase the intelligence of their network operations, planning and optimization.

Heavy Reading believes that artificial intelligence (AI) and machine learning (ML) will be key to automating network operations and enhancing the customer experience. Although "big data" analytics is already widespread in the telecom industry, it is typically conducted in batch, after the fact, and used to manually update rules and policies. In order to move to real-time closed-loop automation, CSPs need systems that are capable of learning autonomously. That is only possible with AI/ML.

Researchers in communication networks are tapping into AI/ML techniques to optimize network architecture, control and management, and to enable more autonomous operations. Meanwhile, practitioners are involved in initiatives such as the Telecom Infra Project's (TIP) Artificial Intelligence and Applied Machine Learning Group. AI/ML techniques are beginning to emerge in the networking domain to address the challenges of virtualization and cloud computing. Network automation platforms such as the Open Networking Automation Platform (ONAP) will need to incorporate AI techniques to deliver efficient, timely and reliable operations.

However, we must not let ourselves get carried away by the breathless hype surrounding AI/ML. Many so-called AI/ML systems today are mainly composed of "big data" tools, statistical analysis and a healthy dose of marketing. As our sister market intelligence firm Tractica surmises in its report Artificial Intelligence for Telecommunications Applications: "An immature ecosystem for telecom AI use cases has formed, made up of legacy telecom network and business support system (BSS)/operations support system (OSS) vendors; broad-based automated customer service specialists; CRM providers; open-source communities and organizations; established cybersecurity vendors; and a small but impressive number of startups."

To help navigate the AI/ML topic, this report provides an overview of AI/ML, outlines the key telecom use cases, quantifies the level of adoption in CSPs today, and discusses the challenges of applying AI/ML to the networking domain. The report also provides real-world examples from 10 CSPs using AI/ML and summarizes key AI initiatives taking place in academia (Knowledge-Defined Networking), standards organizations (ETSI and IEEE), industry consortia (TIP) and open source projects (Acumos). Finally, the report profiles 16 vendors that are specialists in telecom AI/ML.

1.1 Key Findings

The key findings of this report are as follows:

Progress in AI/ML has accelerated in recent years. Drivers include breakthroughs in neural network theory, the availability of massive data sets for academics to experiment with, and the rise of public cloud, making computing capacity readily available and cheap.

The top AI use case in telecom is network operations monitoring and management. Other popular use cases include predictive maintenance; fraud mitigation; cybersecurity;

customer service and marketing virtual digital assistants; intelligent customer relationship management (CRM) systems; and customer experience management (CEM).

Of the 10 CSP case studies profiled in this report, eight of them describe using AI in networking, seven in customer care and two in fraud/security. Of the 16 vendors profiled in this report, nine of them are applying AI to networking, nine to customer care, four to fraud/security and four to marketing/CRM.

In a 2017 Heavy Reading survey, most respondents said that AI/ML would become a critical part of network operations by 2020. In a 2017 survey by TM Forum, 52 percent of respondents claimed they were already using ML and analytics for network management, while another 38 percent planned to do so in the next two years.

The greatest challenge to applying AI/ML to networking – as opposed to multi-industry applications such as CRM or chatbots – is data being unavailable, difficult to access or "dirty." Other challenges include the lack of data science talent; the lack of a clear question to answer; and the limitations of existing tools.

These challenges may be addressed by initiatives currently ongoing in the industry. Examples include the European Telecommunications Standards Institute's (ETSI) Experiential Networked Intelligence Group; the Telecom Infra Project's AI and Applied ML Group; and the Linux Foundation's Acumos open source project.

Although Heavy Reading is optimistic about the potential for AI/ML in telecom, we are also aware that there is a great deal of hype around these terms. Many industry participants use these acronyms inappropriately in their marketing material when, in fact, their systems are still based on a traditional rules-based approach.

Nonetheless, an increased level of automation is required to manage virtualized networks (including 5G), and AI/ML could play an important role, particularly in supporting real-time decision-making. ML and AI promise to reveal new insights from network telemetry and flow data, enabling CSPs to predict capacity demands and scale their networks appropriately.

1.2 Companies Covered

Vendors covered in this report include:

- Afiniti Inc.
- AIBrain Inc.
- Anodot Ltd.
- Arago GmbH
- Aria Networks Ltd.
- Avaamo Inc.
- B.Yond Inc.
- Cardinality Ltd.
- Guavus Inc.
- Intent HQ Ltd.
- IPsoft Inc.
- Nuance Communications Inc.
- Skymind Inc.
- Subtonomy AB
- Tupl Inc.
- Wise Athena Inc.

2. INTRODUCTION

In Brian Levy's article How Will AI and Machine Learning Impact CSPs?, he describes AI as "the capability of a machine to imitate human behavior. Machine learning, which evolved from the study of pattern recognition and computational learning theory in AI, explores construction of algorithms that can learn from and make predictions about data."

In Pedro Domingos' book The Master Algorithm, he writes that "Machine learning is a subfield of AI, but it's grown so large and successful that it now eclipses its proud parent."

As Levy goes on to explain, in traditional programming, a human writes a computer program and provides the data, which the computer processes to create the output. In ML, humans provide the data along with the desired output, rules and constraints, and the computer writes the program to deliver this. The two contrasting approaches are shown schematically in **Figure 1**.

Figure 1: What Is Machine Learning?

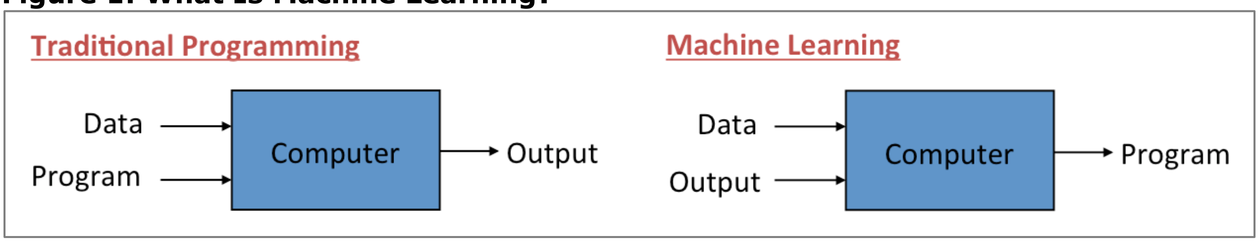

Source: Hydrasky.com

2.1 Machine Learning Categories

ML tasks can be broadly categorized into three main categories: supervised learning, unsupervised learning and reinforcement learning, as shown in **Figure 2**.

Figure 2: Machine Learning Classification

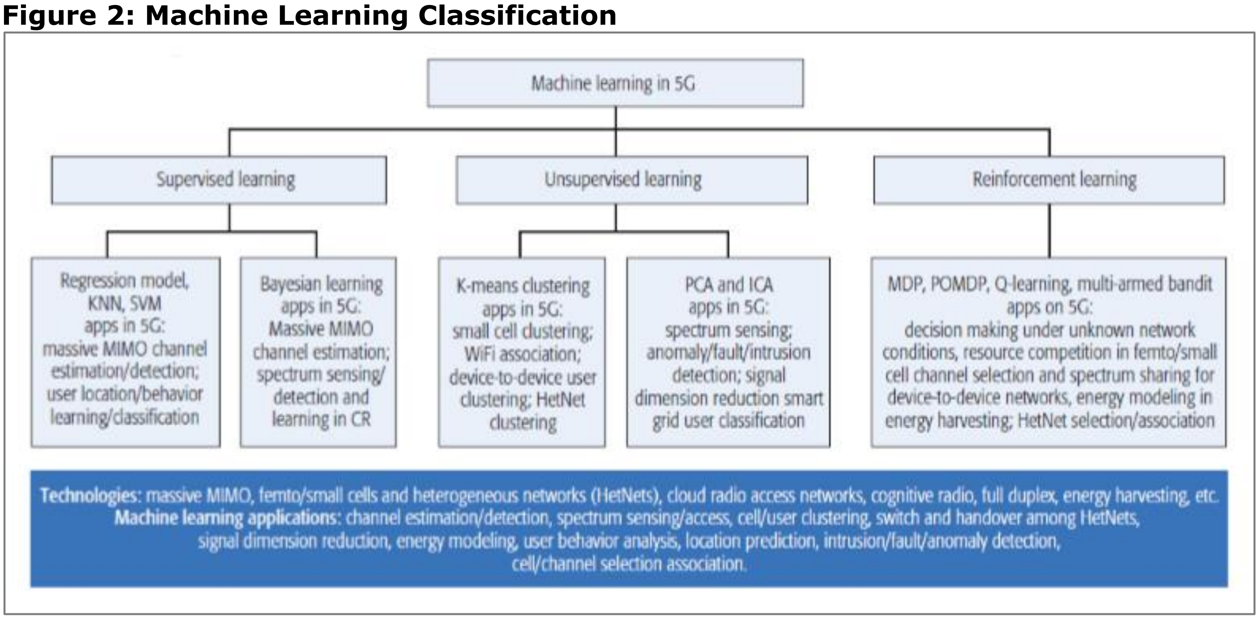

Source: C. Jiang, et al., Machine Learning Paradigms for Next-Generation Wireless Networks

Supervised Learning

With supervised learning in addition to the data inputs we also show the system the desired outputs (e.g., what a human operator would choose to do given a set of input data) and ask the system to create its own mapping of inputs to output conditions (such that it might act autonomously in the future).

Supervised learning can be further broken down into two subcategories:

- Classification systems, e.g., Bayesian networks; Decision trees; Logistic regression; Random Forest; Support Vector Machine. Classification algorithms could be used to group and prioritize alarms.
- Regression systems, e.g., KNN; Linear regression. A regression algorithm indicates a statistical relationship between two, or more, variables (e.g., temperature and noise).

After feeding our chosen algorithm with the labeled data, we produce a learned model, which can then be used on new, unlabeled data to make predictions. A typical example of supervised learning is a spam filter that is given examples of spam and non-spam, and then decides which category new emails most resemble.

Unsupervised Learning

With unsupervised learning the system tries to find patterns in the input data without knowing of any specific output conditions of the data. An example might be the detection of anomalous behavior indicating a security threat. The key difference is that supervised learning uses a labeled training set of data, while unsupervised learning must discover patterns from unlabeled data.

Clustering techniques are normally applied to unsupervised learning systems. Clustering also designates input data (e.g., alarms) into different groups although these groups are not predetermined and hence clustering is generally an unsupervised form of ML. Examples of clustering algorithms include Apriori; distribution-based (Gaussian mixture models); and K-Means.

Clustering might be used for anomaly detection, i.e., spotting events that are outside normal cluster behavior.

Reinforcement Learning

Reinforcement learning is an area of ML, inspired by behaviorist psychology. The system is given a goal and adjusts its behavior to maximize its performance. Feedback is provided in rewards and punishments as the system explores the problem, trying to find the optimal solution. Reinforcement learning differs from standard supervised learning in that correct input/output pairs need not be presented. An example of reinforcement learning is learning to play a game such as chess.

Reinforcement learning is particularly well suited to dynamic environments where one does not just need to make a prediction based on some historical data, but also must adapt to a changing environment that one is trying to control. The system observes the environment, takes actions to maximize performance, makes a new observation of the environment, and then takes further action depending on whether the system is now closer to or further from the desired state.

Deep Learning

While the three techniques described above are learning approaches, deep learning is a class of ML algorithm. Most modern deep learning models are based on an artificial neural network. A deep learning neural network uses a cascade of processing layers, whereby each layer transforms the input data into more abstract representations.

For example, the first layer might recognize the edges in an image, the second layer might identify a face, the third layer a feature on the face such as the nose. The output layer combines those features to make predictions such as that the image is the face of a person, or a specific person whose likeness is known to the system. Deep learning has been used for image recognition (computer vision), speech recognition (audio to text) and natural language processing (meaning extraction from audio/text).

2.2 Why the Resurgence of Interest in AI/ML?

AI and ML are not new topics. Even in the context of telecom, as far back as 1993 researchers were exploring AI techniques that they thought would be essential to the transformation of the telecom network. They key factors that have led to an acceleration in progress in ML in recent years include:

- Breakthroughs in neural network theory around 2006
- Improvements in computing capacity: x86 CPUs, GPUs, FPGAs and custom ASICs designed specifically for ML, e.g., Google's Tensor Processing Unit (TPU) and associated Tensor Flow software libraries
- Public cloud (AWS, Azure, etc.) making computing capacity highly available and cheap
- Massive data sets: online photos, email, video, gaming, search, messaging, mapping and shopping are fertile ground for ML
- Success stories: AlphaGo, Google Pixel Buds, image recognition, lip reading, etc.

AI has seen renewed interest in the telecom sector, due especially to the rising complexity of network technology. As Mirza Golam Kibria, et al., outline in their November 2017 paper, Big Data Analytics, Machine Learning and Artificial Intelligence in Next-Generation Wireless Networks, "The next-generation wireless networks are evolving into very complex systems because of the very diversified service requirements, heterogeneity in applications, devices, and networks. A novel paradigm of proactive, self-aware, self-adaptive and predictive networking is much needed."

3. POTENTIAL AI/ML USE CASES IN TELECOM

Tractica's report Artificial Intelligence for Telecommunications Applications identifies seven key telecom AI use cases:

1. Network operations monitoring and management
2. Predictive maintenance
3. Fraud mitigation
4. Cybersecurity
5. Customer service and marketing virtual digital assistants
6. Intelligent CRM systems
7. CEM

Below we discuss each of these areas in turn.

3.1 Network Operations Monitoring & Management

AI and ML approaches are beginning to emerge in the networking domain to address the challenges of virtualization and cloud computing. Increased complexity in networking and networked applications is driving the need for increased network automation and agility. Network automation platforms such as ONAP should incorporate AI techniques to deliver efficient, timely and reliable management operations. Examples of network-centric applications of AI/ML include:

- Anomaly detection for operations, administration, maintenance and provisioning (OAM&P)
- Performance monitoring and optimization
- Alert/alarm suppression
- Trouble ticket action recommendations.
- Automated resolution of trouble tickets (self-healing)
- Prediction of network faults
- Network capacity planning (congestion prediction)

ML could support network operations to detect issues – e.g., faults, service-level agreement (SLA) breaches – in real time, diagnose root causes, correlate across multiple event sources, filtering out noise (false alarms), and recommend solutions. 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.

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 ML we could instead train a system to devise its own rules based on a given set of data inputs (e.g., network telemetry).

ML could be applied to service assurance to automate the resolution of common incidents. The system could be taught by operations staff how to handle these common incidents but still require human approval before taking action (supervised or open-loop operation mode). Longer term, as humans become more comfortable with the ML technology, they may let it operate with increasing autonomy.

Professor Maziar Nekovee, University of Sussex, has identified the following potential use cases with AI and ML algorithms in a mobile context:

- **AI at the RAN:** intelligent initial access and handover; dynamic scheduling; resource optimization
- **AI at the core:** autonomous VNF scale in/out, up/down; provision of elasticity; intelligent network slicing management; service prioritization and resource sharing; intelligent fault localization and prediction
- **AI at the fronthaul:** traffic pattern estimation and prediction; flexible functional split
- **Other general AI applications (RAN, core or end-to-end network):** energy efficiency according to dynamic traffic pattern, etc.; end-to-end service orchestration and assurance (e.g., customized SLA); end-to-end service optimization, prioritization

3.2 Predictive Maintenance

Heavy Reading sees predictive maintenance as a subcategory within network operations rather than a separate field. That said, in a 2017 survey of CSPs (see **Figure 3**) we found that predictive maintenance was the top use case for ML in telecom, ahead of security, network management, and fraud/revenue assurance.

Figure 3: Use Cases for ML in Telecom

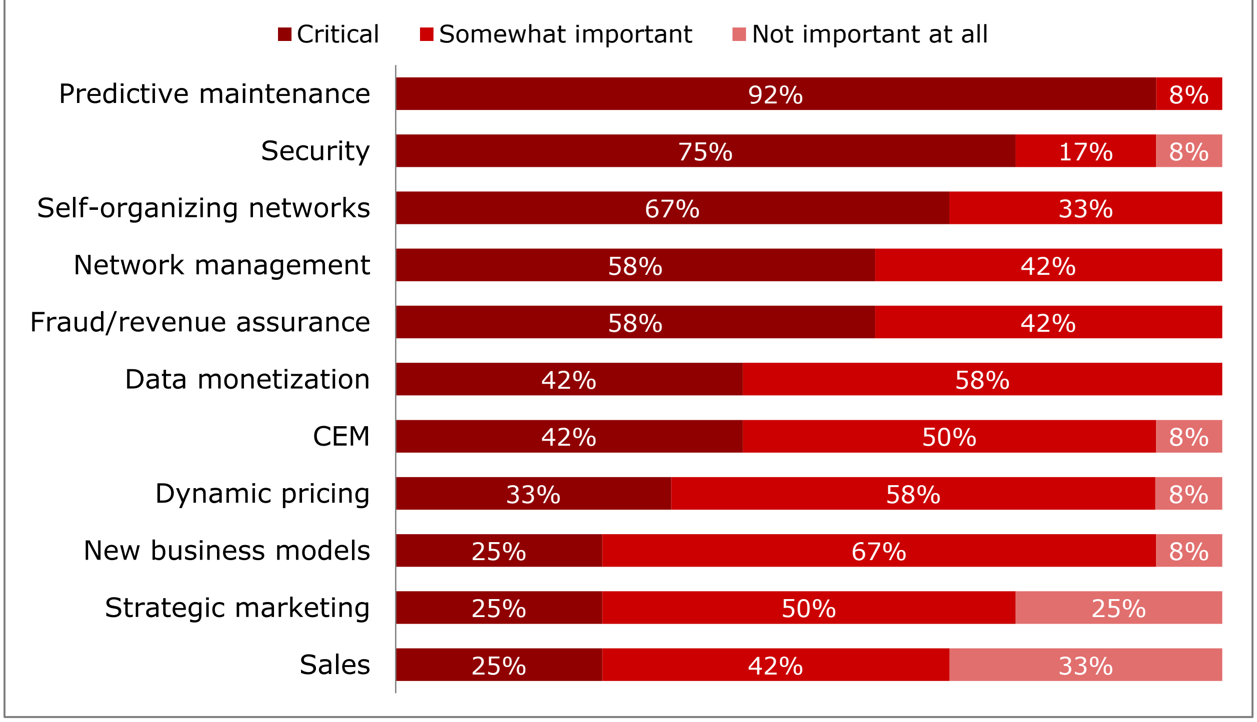

Source: Heavy Reading Survey – Thought Leadership Council (n=12), Nov. 2017

3.3 Fraud Mitigation

Fraud detection and prevention was the fifth most popular use case in the survey results shown above. According to the Communications Fraud Control Association, fraud costs the global telecom industry $38 billion annually, of which roaming fraud accounts for $10.8 billion. In **Section 8**, we describe how:

- Anodot uses AI to identify revenue leakage and surface discrepancies between expected results and how events are actually billed.
- Skymind is using AI to combat subscriber identity module (SIM) box fraud at Orange.
- Wise Athena has used AI to identify CSP fraud.

3.4 Cybersecurity

Security was the second most popular use case in our survey. Heavy Reading's Telecom Security Market Tracker has found that there is guarded optimism over AI for the automation of CSP security.

Traditional security technologies rely on rules and signatures to find threats but this information can soon become out of date. The tactics of adversaries are evolving rapidly, and the number of advanced and unknown threats targeting CSP networks continues to increase. AI/ML algorithms could be trained to adapt to the changing threat landscape, making independent decisions about whether an anomaly is malicious or providing context to assist human experts.

According to our Telecom Security Market Tracker, AI techniques such as neural networks and ML have already been used for many years to improve the detection of malicious code and other threats within telecom traffic. And AI has the potential to go further, such as automatically taking remediation actions or presenting a human security analyst with the right type of data on which to base a decision, and perhaps a recommendation.

One recent hot area of activity is in baselining of the behavior of devices connected to the Internet of Things (IoT). Here many established vendors and AI startups are developing solutions that will help CSPs to manage IoT devices and services more securely, making use of automatic profiling of those devices.

Network security (including DDoS, threat discovery and mitigation) is cited as one of the applications for Colt's project Sentio. In **Section 8**, we describe how Arago and Guavus are used for security operations.

3.5 Customer Service & Marketing Virtual Digital Assistants

One of the key applications of AI/ML in the telecom sector to date has been the use of chatbots to augment or replace human call center agents. Seven of the 10 CSP profiles in this report include a discussion of their use of AI in customer care.

For example, Telstra's Kieran O'Meara, Director Technology Design & Delivery, estimates that 30 percent of inbound calls to a contact center could be resolved by AI chatbots. There is still a role for human agents at Telstra (it has 11,000 today), but with AI assistance O'Meara estimates that they can be 20-35 percent more productive. Telstra has around 300 agents managing chatbots on its websites but doesn't expect this number to grow. Instead, it plans to increase the number of agents dealing with customer enquiries directly via messaging apps such as WhatsApp.

Other examples of AI usage in customer service/support include:

- Knowledge portals and AI assistants for human agents
- Contact center optimization and compliance
- Customer voice and text sentiment analysis – Telstra is looking at using text sentiment analysis to enhance the performance of its messaging and chat agents.

3.6 Intelligent CRM Systems

AI can be applied to CRM in areas such as personalized promotions, cross-sell/up-sell opportunity identification, and churn prediction and mitigation. Research by a group called Wise Athena (see **Section 8.14**) investigated the use of deep learning to predict customer churn in a mobile telecom operator. They found the method more accurate than previous methods based on supervised ML classifiers. Other vendors we profiled that are using AI to provide marketing insights include Cardinality, Guavus, and Intent HQ.

3.7 CEM

Heavy Reading sees Customer Experience Management as the process of managing "all customer touchpoints" to ensure a positive relationship with the brand. As digital touchpoints continue to grow, analytics and AI are essential tools for CSPs to understand the health of the network, the customer journey (customer care, billing, etc.), and real-time service quality. As such the CEM category intersects customer service, marketing, CRM and the service assurance side of network operations and management.

4. CSP ADOPTION OF AI

Below we discuss the findings of surveys by TCS, TM Forum and Heavy Reading that shed light on the degree of adoption of AI by CSPs.

4.1 TCS Study Suggests Most CSPs Already Using AI/ML in IT/Networking

A 2017 survey by TCS, Getting Smarter by the Sector, of 59 CSP representatives across North America (39 percent), Europe (29 percent), Asia/Pacific (22 percent) and Central/Latin America (10 percent) found that 93 percent of their employers were already using AI while the remainder planned to do so by 2020. By far, the largest use for the technology is in IT (which here includes networking) including tasks such as resolving internal users' tech problems, detecting and deterring security intrusions, ensuring that all vendors are from approved vendor lists, and automating IT production management. Telecom companies also commonly use cognitive tools for customer service, sales, and marketing work.

Figure 4: Telecom Companies Using AI per Function

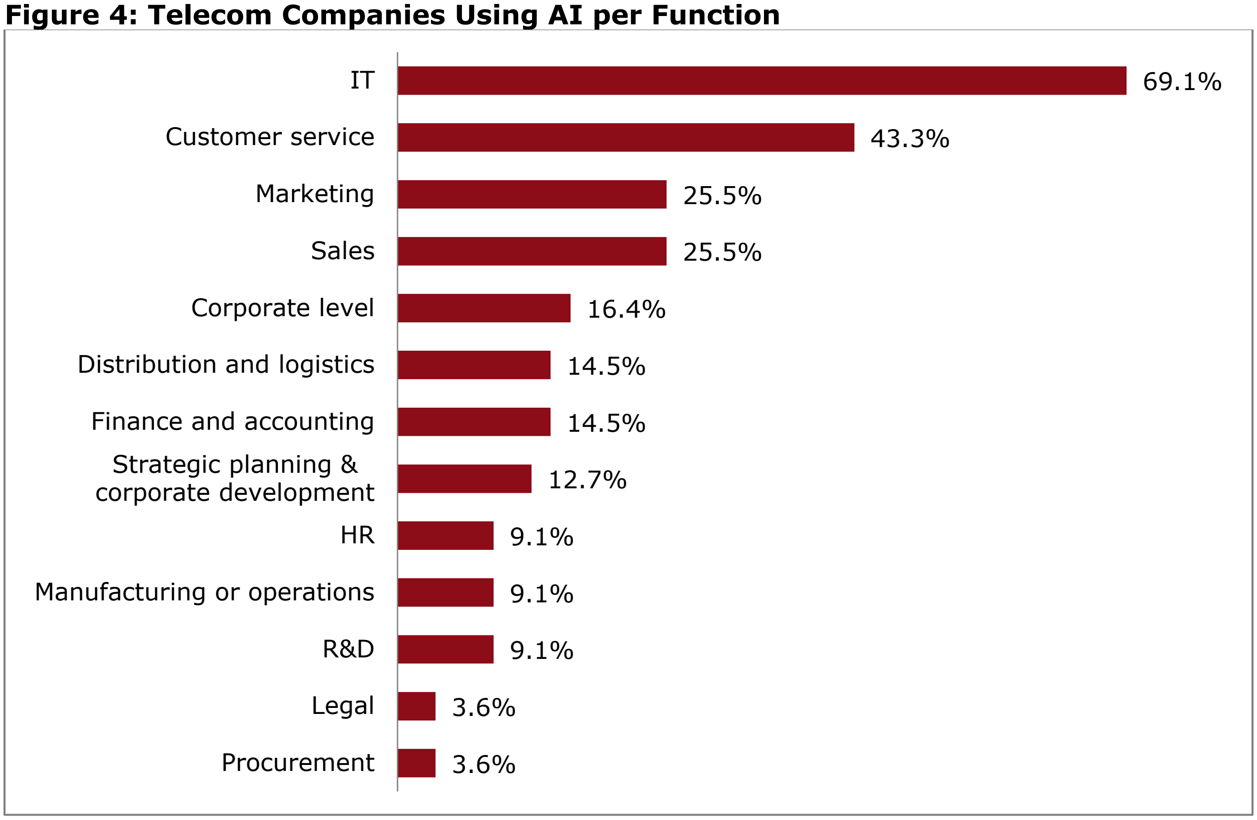

Source: TCS, 2017

4.2 TM Forum Study Suggests More Cautious Adoption of AI/ML

The high AI adoption rate of the TCS study contrasts with a 2017 TM Forum survey of 187 executives from 76 CSPs operating in 51 countries, which found that only 18 percent were incorporating AI into their product and service roadmaps. Clearly the survey asks a different question, and the demographics may differ; nonetheless, the difference in results is striking. Note the "wait and see" camp in the TM Forum survey were mainly from developing countries, or from operating companies within large groups where AI projects are managed centrally.

Figure 5: What Is Your Organization's Response to AI?

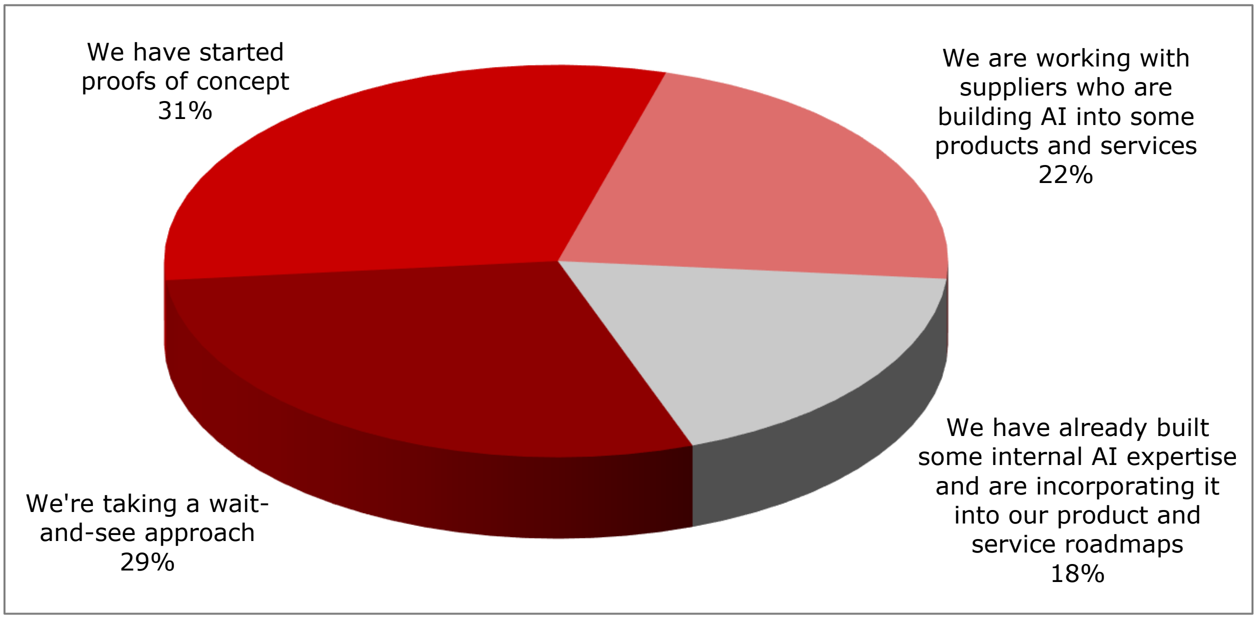

Source: TM Forum, 2017

4.3 Customer Experience the Key Driver for AI/ML

Given a choice of three options, the TM Forum respondents indicated that customer experience was the key driver for AI. To that point, 30 percent of CSP respondents said that they have rolled out chatbots; 71 percent of these have retained human agents who are ready to step in to deal with complex interactions. One in seven have redeployed some customer service agents to complete higher-value tasks.

Figure 6: What Is the Main Driver for AI in Your Company?

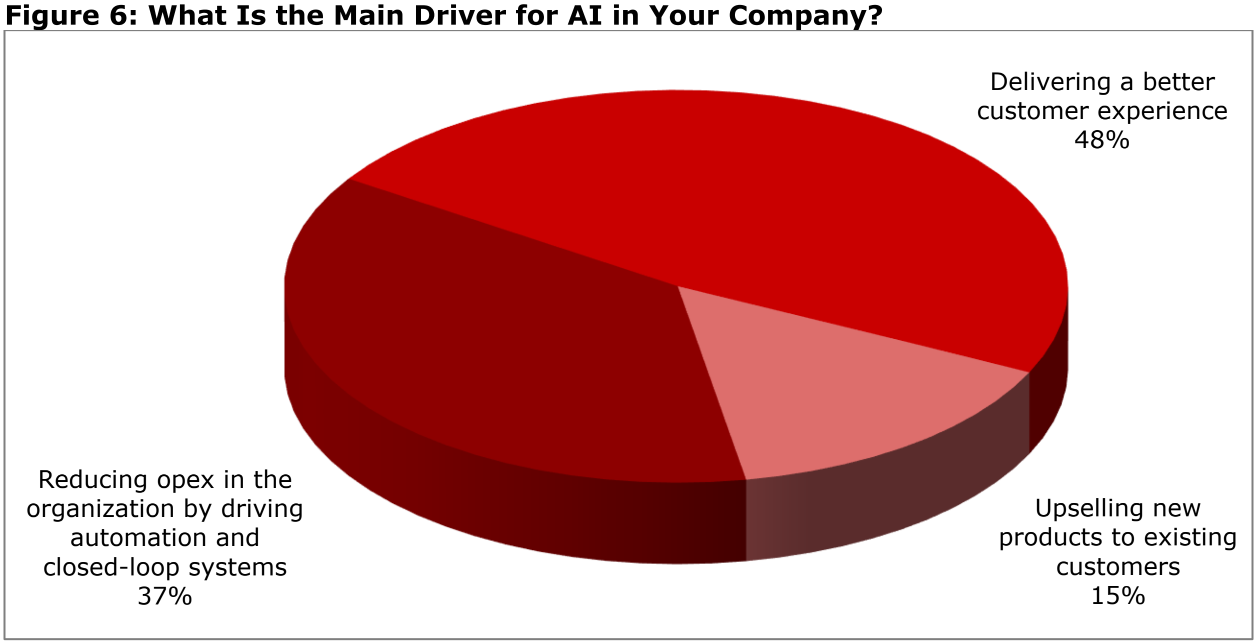

Source: TM Forum, 2017

4.4 AI/ML in Network Management

The 2017 survey by TM Forum (AI: The time is now) found that 52 percent of respondents claimed that they were already using ML and analytics for network management, and a further 38 percent planned to do so in the next two years. Examples of how they were using ML and analytics included problem prediction, self-healing and automated policy change.

Figure 7: How Are You Using AI/ML Plus Analytics for Network Management?

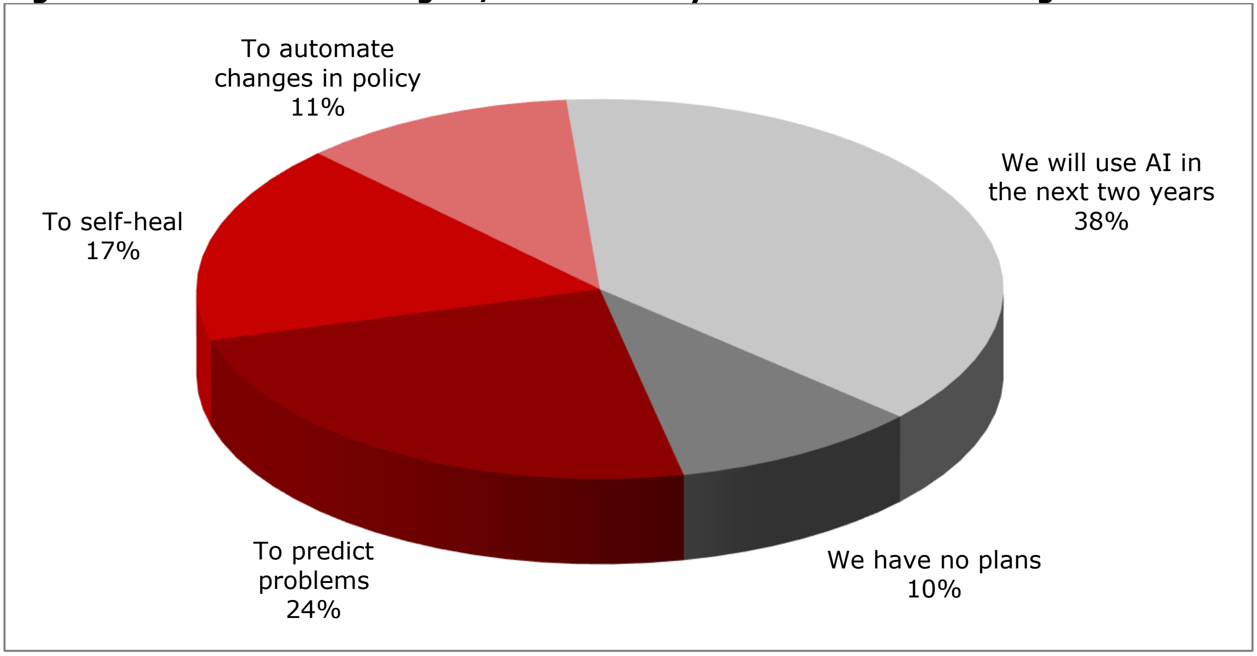

Source: TM Forum, 2017

In a Heavy reading 2017 survey, most respondents said that ML would become a critical part of network operations by 2020.

Figure 8: When Will ML Become a Critical Part of Your Company's Network Operations?

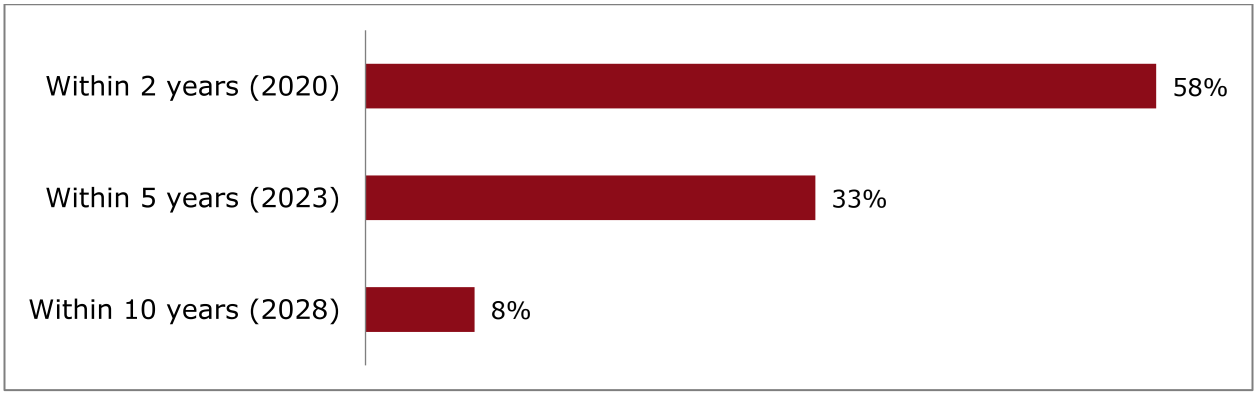

Source: Heavy Reading Survey – Thought Leadership Council (n=12), Nov. 2017

5. REAL-WORLD CSP EXAMPLES

Although there is great interest in the topic of ML, there are few public examples of how CSPs are using it in their operations today. Below we reference some real-world examples in customer care, networking, and security. **Figure 9** is not an exhaustive list of all uses cases of these companies, or indeed of all CSPs using AI.

Figure 9: CSP Example Summary

Company	Customer Care	Networking & IT Ops	Fraud & Security
AT&T	✓	✓	✓
Colt	✓	✓	
Deutsche Telekom	✓		
Globe Telecom	✓	✓	✓
KDDI		✓	
KT		✓	
SK Telecom		✓	
Swisscom	✓		
Telefónica	✓	✓	
Vodafone	✓	✓	

5.1 AT&T

According to AT&T's Inside Connections Blog, the company has been building AI and ML systems for decades, using algorithms to automate operations such as common call center procedures, technician dispatching, and to analyze and correct network outages. AT&T says it is now using AI to help make its networks more secure, self-healing and self-resilient as it prepares for the rapid growth of video traffic, particularly on mobile networks.

AT&T Labs Research has described how it is using ML to create a "virtual world" that describes its 5G infrastructure and environment – poles, buildings, building materials, foliage – to help determine where cell sites can be placed without requiring a site visit and identify faults in towers. AT&T also sees AI as key to enabling organizations to take advantage of the IoT.

AT&T has published a book called Artificial Intelligence for Autonomous Networks (out this October), which explores the potential to transform network operations, cyber security, enterprise services, 5G and IoT, infrastructure monitoring, traffic optimization, customer experience and care.

The book's editor published a blog late last year (Pay No Attention to the Man Behind the Curtain: A Reality Check for Artificial Intelligence) warning that "many of the things we think of as AI today in fact require people to manually input and structure massive amounts of data. The end user might be interacting with software, but behind the scenes, there are a lot of people sweating to make that software look smarter than it really is. For example, even commercial AI solutions that are used today, such as virtual assistants and call routing systems, are developed on an enormous amount of data that is labeled manually. Extensive human effort goes into making one AI application a reality."

AT&T launched the Acumos project (see **Section 7**) as a way of reducing some of the effort required to build AI applications.

5.2 COLT

Colt has created a new "AI-driven networking" project called Sentio with the aim of developing fully automated service management capabilities. The Sentio project started last year and Colt is currently implementing a proof of concept (PoC).

Colt cites the following applications for Sentio:

- Traffic flow classification
- WAN path optimization
- Fault prediction
- Quality-of-experience (QoE) modelling
- Intelligent bandwidth on-demand
- Capacity management
- Network scaling
- Network security – DDoS, threat discovery and mitigation
- Service modification and restoration through the automated scaling of VNFs.
- Network operations automation – Colt aims to enhance its service assurance capabilities by taking non-traditional data (signal strength, power, temperature, etc.) from network elements (cards, links, etc.) to predict potential faults.
- Customer experience – helping customer service teams to deal with customer inquiries through chatbots and quicker access to relevant information (e.g., known faults).

5.3 Deutsche Telekom

Instead of buying "off-the-shelf" AI systems and robots, which can be expensive, Deutsche Telekom is developing its own AI solutions – via its own developer teams and partners. One area that is ripe for automation with AI is resolving queries from enterprise customers, which Deutsche Telekom notes can sometimes require 1,000 manual actions in various software systems (entering bookings and process commands, initiating orders, etc.).

Deutsche Telekom already uses chatbots or digital assistants to relieve human agents of standard tasks. Virtual assistants such as Tinka, Sophie and Vanda will soon be able to "learn" from chat logs and from real conversations between service agents and customers. They are also being designed to communicate with customers in ways that sound and feel "human."

- Tinka is a chatbot for the Austrian market that has learned more than 1,500 answers so far. Tinka is able to handle about 80 percent of all questions put to her. When she can't answer a question, she forwards it to a human. Tinka can support customers in setting up LTE-based home WiFi networks and inserting SIM cards into phones.

- Sophie is a chatbot for Deutsche Telekom's no-frills youth brand, Congstar that works in conjunction with the Facebook Messenger application.
- Vanda is a natural language processing chatbot used for enterprise self-service in Deutsche Telekom's eastern European markets (e.g., Hungary) via Facebook, Viber and other messaging and voice platforms.

Deutsche Telekom has launched an overarching AI program, eLIZA, for the purpose of linking all AI solutions within the group. Other examples of how Deutsche Telekom is using AI can be found here and here. Deutsche Telekom is also active in the AI and Applied ML Project Group launched by TIP (see **Section 7.3**).

5.4 Globe Telecom

Globe Telecom has been experimenting with AI in customer-facing scenarios as part of the TM Forum Catalyst Program and is one of the most vocal proponents of the TM Forum working group. Vincent Seet, Head of Enterprise Architecture, explained that at Globe AI use cases typically can be categorized using an "80-20 rule," which means that 80 percent of the time they can be developed on top of and be supported by a centralized platform (e.g., chatbots, fraud management). The rest of the time, they need to be tightly coupled point solutions (e.g., self-organizing networks).

Figure 10: Most AI Use Cases Should Be Managed With a Common Platform

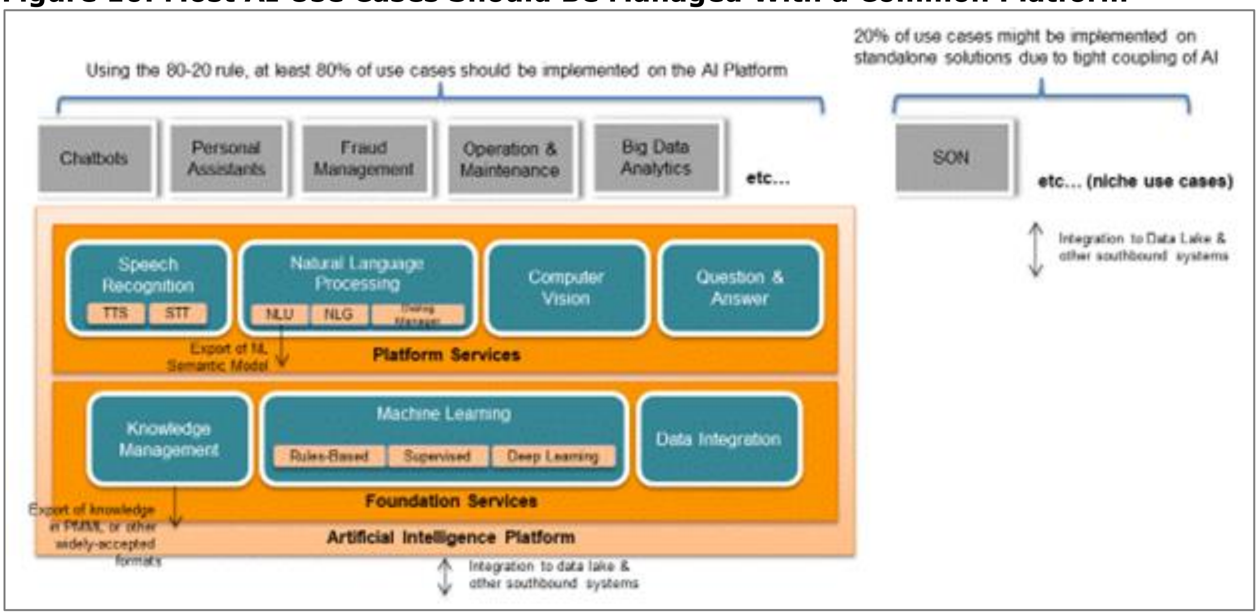

Source: Globe Telecom, TM Forum

5.5 KDDI

In 2016 the R&D Labs of Japanese operator KDDI announced an AI-based monitoring system that was used to predict anomalies caused by NFV hardware and software. The AI-based monitor learns what are normal and abnormal the conditions, and the software-defined networking (SDN)/NFV orchestrator provides an appropriate recovery plan so that affected services are quickly preserved. KDDI described the PoC as a world first and an important step toward the network virtualization era supporting 5G.

Figure 11: KDDI's AI-Assisted Automated Operation Systems

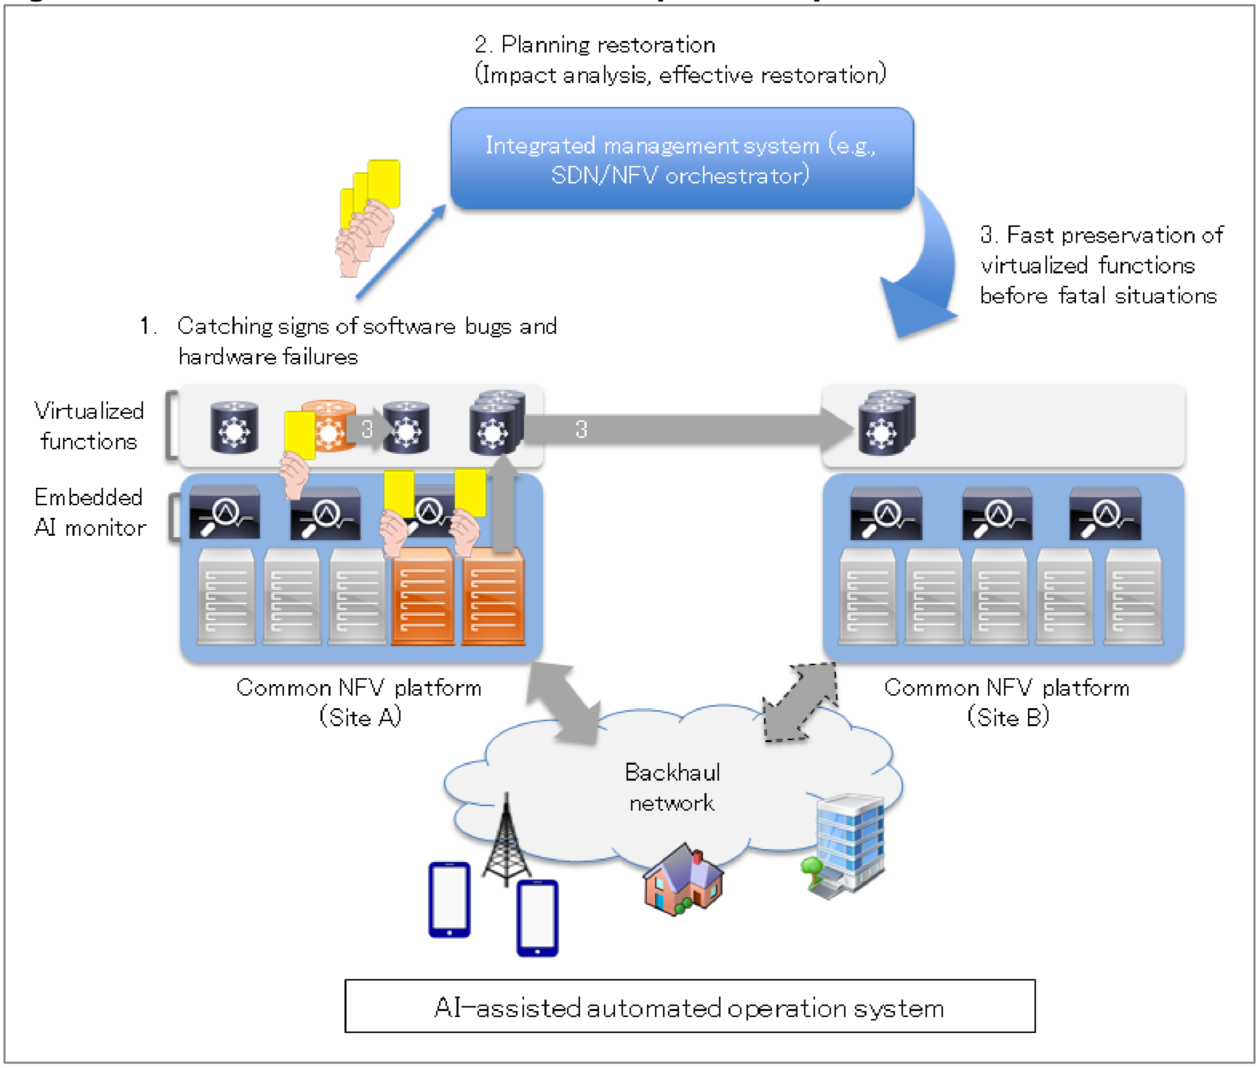

Source: KDDI

5.6 KT

Earlier this year, South Korean CSP KT announced the development of a new AI platform, Neuroflow, for network operations. According to [Korea Times](), KT plans to establish an AI-assisted network control center by 2020. Neuroflow has been made open source, allowing other companies to use it. A search for "Neuroflow" on the KT website yielded no results.

5.7 SK Telecom

In October 2017, SK Telecom announced the expansion of its T Advanced Next Generation Operational Supporting System (TANGO) from the fixed to the mobile side of its business. TANGO is an AI-assisted network operation system with big data analytics and ML capabilities that detects issues on the network, troubleshoots problems, and the optimizes performance. SK Telecom has also signed a strategic partnership with Indian operator Bharti Airtel for the use of TANGO.

SK Telecom has another AI initiative called AIX (AI Inference Accelerator) which relates to the data center infrastructure of the company's "Nugu" AI speaker device. AIX is a palm size, card-shaped accelerator that can be installed to the existing AI server inside Nugu's data center infrastructure to enable 20 times faster computing of deep-learning technology.

5.8 Swisscom

Swisscom is introducing its own AI chatbots for customer services. Pascal Jaggi, the head of Customer Care at Swisscom, has presented the operator's home-grown AI system Cosmos, which can distribute written enquiries from customers to the right employees. Swisscom's service center processes 30,000 to 40,000 written enquiries per month. Cosmos can process and forward enquiries with an 85 percent reliability, far greater than typical human results. Cosmos can also reply automatically, but currently this is limited to only a small fraction of enquiries, which can be redirected to self-care.

Cosmos is based on a deep learning approach and was "fed" with 800,000 written requests prior to launch. In addition, 100 employees are currently investing two hours each day to further train the AI. Swisscom is also working on a chatbot that will be connected with Cosmos.

Swisscom has a similar initiative for its enterprise customer support. Marmo is deployed in corporate business and uses AI to automatically seek solutions across all sources in the existing infrastructure. The call agent enters a couple of keywords and the system automatically checks for similar requests in the past and offers a possible solution in a matter of seconds. The call agent then decides whether the proposed solution is useful.

5.9 Telefónica

At Mobile World Congress this year Telefónica presented its "fourth platform," originally announced a year earlier. The first three platforms (physical assets, IT systems and products/services) are table stakes for CSPs; Telefónica's fourth platform concept is akin to the "digital service provider" model other CSPs are striving toward.

The "fourth platform" uses AI to analyze data from the underlying platforms to better serve customers. Specifically, Telefónica has launched its voice-activated "cognitive" assistant, Aura, in six markets – Argentina, Brazil, Chile, Germany, Spain and the U.K. Telefónica says that Aura will learn from its interactions with individual customers and ultimately be able to provide tailored recommendations and support based on a user's preferences.

Telefónica is not only applying big data analytics and ML to its customer-facing activities but also to improve the operational efficiency of the business. In 2007 Telefónica first started applying algorithms and ML to troubleshooting in its network operations centers (NOCs). The team of data scientists from Telefónica R&D that developed various tools for NOCs have published several academic papers and filed multiple patents related to their work. Leveraging this small group of data scientists is a team of hundreds that apply business intelligence and data visualization tools to operational and commercial use cases across Telefónica.

The first part of the process is creating the data repositories themselves. Telefónica has a separate data lake in each operating business (Spain, Germany, U.K., Brazil, etc.) and a centralized, global big data platform for analysis across the group. In total Telefónica collects data from more than 170 sources of information, including contact center calls, field technician reports, bills, energy usage, OSS, network telemetry and so on. Over time, more and more data sources are being added.

Once data has been collected and anonymized, it must be normalized using a standard data model and checked for quality. Poor data from inventory systems has required Telefónica to replace some systems (e.g., transport and access inventory) and to change some operational processes to improve the quality of the data that is entered into these systems. Even if a perfect network inventory is created, it will soon diverge from reality if field technicians fail to report changes such as a change in port on a device to resolve a trouble ticket.

Once a reliable data set is available, the focus shifts to analyzing use cases. These normally come from business units looking for solutions to real-world problems. So far, Telefónica's data analytics team has worked on solutions for around 300 use cases, the benefits of which it tracks on an ongoing basis. Most of the use cases are operations related (e.g., infrastructure management, customer experience, customer service delivery, internal plant management, etc.). The next largest category is technology-related (video platform, radio planning, etc.).

Figure 12: Telefónica AI & Analytics Use Cases

Use Case	Description
Base station profitability	Total cost is based on rental (data coming from real estate team), maintenance (data from operations), field technician costs, and necessary level 3 support in the NOC. Traffic and associated revenue is derived from commercial team. The profitability of each base station is calculated and an assessment is made of the least profitable base stations to understand what can be changed, e.g., relocating base station 50 m to a different site with lower rental cost.
Preventive maintenance	Normal practice is to replace components periodically based on the vendor's recommended schedule. By collecting its own history of equipment faults, Telefónica can make more accurate predictions of faults based on the specifics of its cell sites (e.g., ambient temperature and humidity). So far, this has led to the reduction of hundreds of site visits in one operating business (country).
Battery capex optimization	Most of the batteries (backup power supplies) deployed in the field are never actually used, though their theft and replacement represents a significant cost. Telefónica has analyzed which sites have historically suffered from low electrical supply reliability and is now focusing the replacement of stolen batteries where the probability of them being needed is highest.
Trouble ticket prioritization	Traditionally trouble tickets are prioritized based on the estimated number of customers impacted and the length of time the ticket has been open. Now when a ticket is opened Telefónica is predicting how long the ticket is likely to take to resolve and is prioritizing tickets based on this measure rather than the time elapsed so far.

Source: Telefónica

Although many of the analytics use cases that Telefónica has developed are based on traditional statistical techniques, some of them also incorporate AI algorithms. For example, Telefónica is currently exploring how to make UNICA (Telefónica's end-to-end network virtualization project) resources more intelligent by using AI. They are also exploring AI to suggest next best actions for staff in the service operations center to resolve issues more quickly. These techniques are equally relevant in orchestration. In video operations, Telefónica is using AI to detect anomalies and transfer customers onto a different headend before their service is impacted.

Another use case where the analytics team has employed AI is to create a real-time index of customer satisfaction. Just measuring traditional key performance indicators (KPIs), such as dropped calls and throughput, does not always correlate well with customer experience (as determined by survey data), especially for complex services such as VoLTE and IPTV. As such, Telefónica has turned to some sophisticated ML algorithms that use network KPIs collected every 15 minutes to predict the customer's satisfaction level with an accuracy of around 60 percent (target 75 percent by year end).

5.10 Vodafone

Vodafone introduced the first live chatbot in the U.K. telecom market, TOBi, in 2017. Vodafone claims the AI-enabled bot, based on IBM's Watson technology, provides relevant support to resolve more than 70 percent of customer queries. Vodafone is also trialing a voice biometrics system and a voice assistant that will be compatible with Amazon's Alexa software and Echo speaker.

On the network side, in 2017 Vodafone announced a trial of ML in a centralized self-organizing network (C-SON), to identify the optimal settings to deliver VoLTE services across and to predict locations where 3G traffic will peak in the following hour. The predictions enable the network to self-configure itself automatically to balance the traffic load among neighboring cells and improve the customer experience. Initial results confirmed an average 6 percent improvement in the mobile download speed and lower interference at the cell sites (the cause of dropped calls, problems connecting and higher device battery drain).

6. CHALLENGES OF APPLYING AI/ML TO NETWORKING

Generic AI applications such as marketing and chatbots apply to many industries, not just telecom. When it comes to applying AI/ML to networking, the industry is at a more nascent stage, partly because these problems are specific to communications, and partly because there is a lack of public data sets on which academics have been able to experiment.

In the paper Big Data Analytics, Machine Learning and Artificial Intelligence in Next-Generation Wireless Networks, the authors highlight the challenges of adopting big data analytics and AI in the next-generation communication system:

- The process of managing and leveraging of a huge amount of data, designing algorithms for dynamic and effective processing of sizable data sets and then exploiting the insights from the data analytics in networks can pose unique challenges.
- The prime concerns for the MNOs emerge from the extent of effort, skills, and workforce needed to manage and operate a big data platform.
- However, the most important and difficult challenge is more likely to stem from the loss of direct control that the MNOs still have over the wireless network. The loss of direct control is incurred from the combination of automation and real-time operations within the big data analytics framework.
- On top of these, a substantial investment is necessary.

Below we detail several barriers to applying data science and ML in telecom (see **Figure 13**). Some of these are generic to any industry, but some are more specific to telecom.

Figure 13: Barriers to Data Science/ML in Telecom

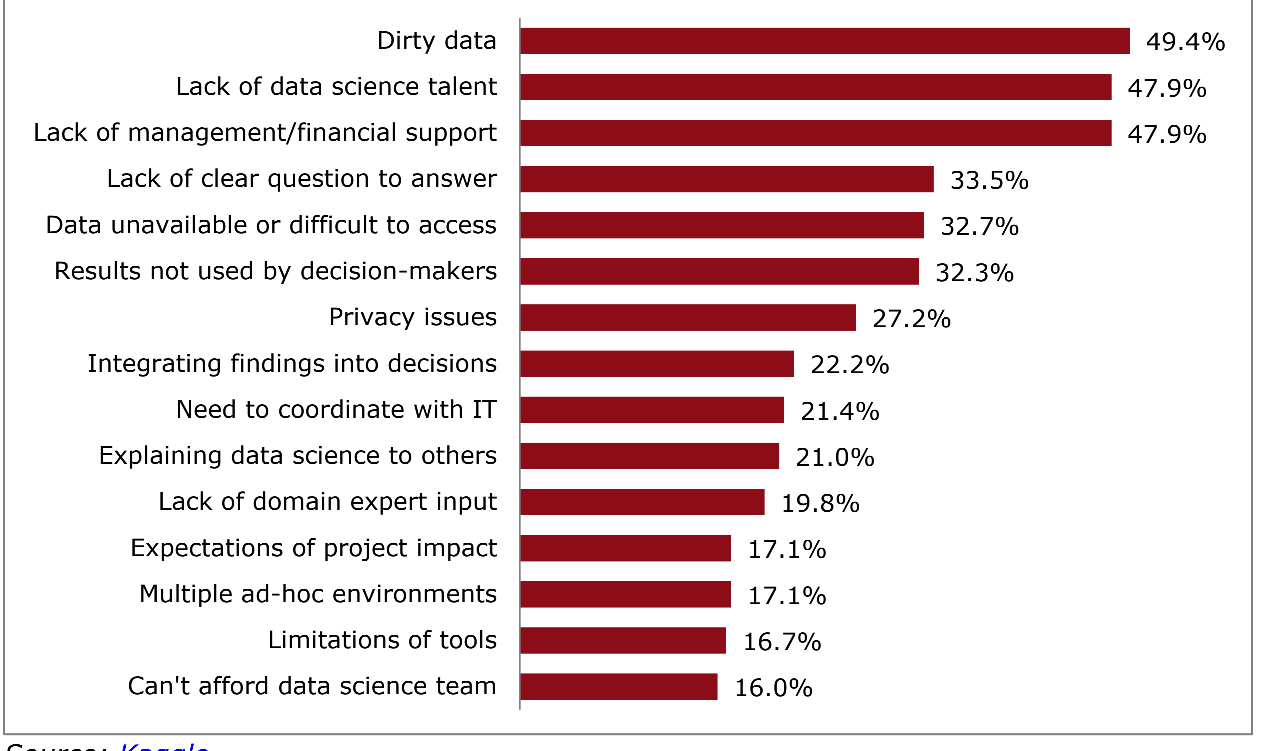

Source: Kaggle

6.1 Data That Is Dirty, Unavailable or Difficult to Access

Telecom networks generate significant amounts of data every day and this is likely to increase further with the move to NFV, 5G and IoT. This data may be stored in a "data lake" but will require cleaning and categorizing before it can be used to train an ML system.

There are diverse types of network data, including flows, logs and various KPIs with no obvious or consistent way to combine them. There is no standard way to integrate network data with other data sources such as CPU and memory usage. Data sets are often noisy, incomplete and not correctly normalized or labeled.

Standardized data sets have been a crucial factor in the success of ML. Examples include:

- The Modified National Institute of Standards and Technology (MNIST) database of handwritten digits
- The ImageNet database
- The Allen Institute for Artificial Intelligence's (AI2) extensive database of datasets

Standardized data sets enable direct comparison of learning and inference algorithms. The result has been the steady ratcheting down of error rates on perceptual tasks (e.g., voice and image recognition) to superhuman levels.

However, in his presentation Machine Intelligence and Networking: Challenges, Opportunities and Realities, David Meyer, Chief Scientist at Huawei Future Network Theory Laboratory, argues that "We have nothing like this for networking:

- "Many/most data sets are toy, noisy, unnormalized
- "Many data sets are proprietary
- "Much data [is] non-iid [independent and identically distributed random variables]
- "Most of our data sources (e.g., NetFlow) are representations of network data that were *not* built for ML."

6.2 Lack of Data Science Talent

Network engineers typically don't have backgrounds that include the kinds of mathematical training and experience that are essential in ML. Recruiting people with the right skills is a challenge. David Mayer also highlights the shortage of AI talent in the networking space: "Network engineers (us!) typically don't have backgrounds that include the kinds of mathematical training and experience that are essential in the ML space, e.g., data modeling and evaluation, software engineering and system design, ML algorithms and libraries."

The result is that there is a serious skills gap. Open source frameworks, such as Acumos, are easing this problem. However, deploying ML at scale still requires intimate understanding of the mathematics, which is a scarce resource today inside CSPs.

As **Figure 14** indicates, the introduction of AI is expected to lead to a change in the skill sets of employees in most organizations, not just telecom. However, at a personal level people are conflicted, hoping on the one hand that AI will take away some of the mundane aspects of their job, but also fearing that there may not a job for them in the long run.

Figure 14: How Do You Expect AI to Affect the Workforce in the Next Five Years?

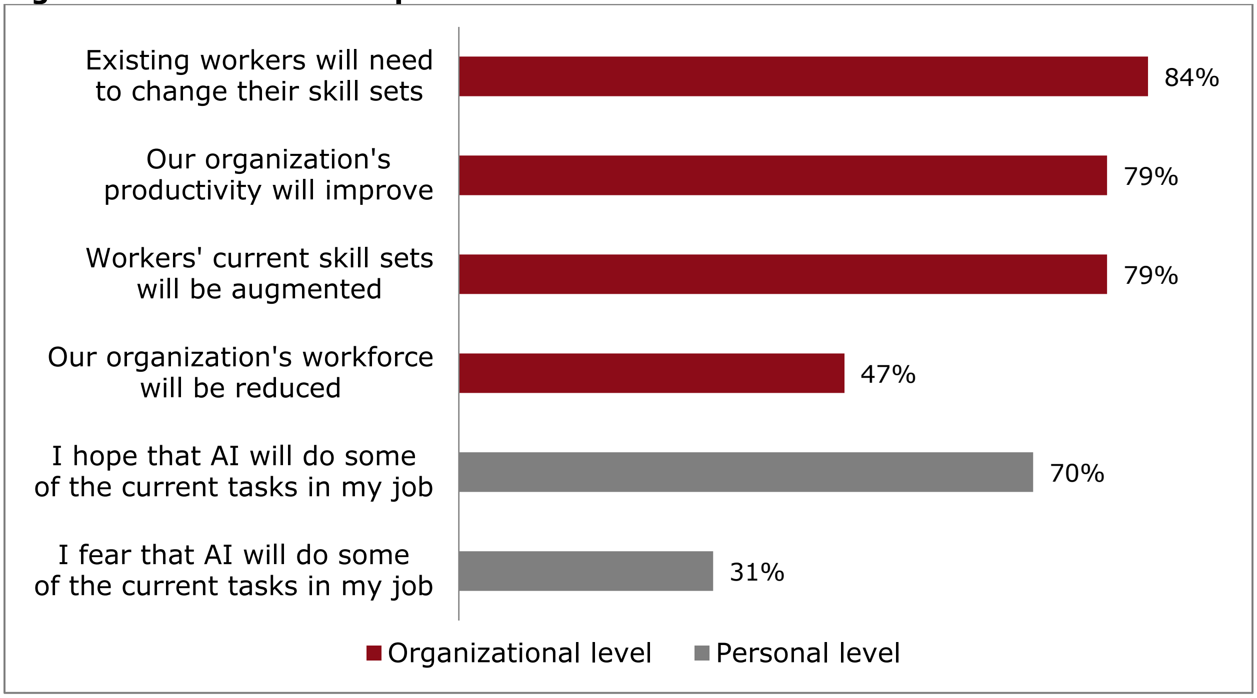

Source: Joint Boston Consulting Group-MIT Sloan Management Review survey on the impact of AI on business, 2017
Note: Shows the percentage of respondents who "somewhat agree" or "strongly agree" with each statement.

6.3 Lack of a Clear Question to Answer

Unlike image and voice recognition, there is no "Theory of Networking" around which to build an AI optimization algorithm.

As analyst Martin Geddes notes in Fundamentals of Network Performance Engineering, "Networking has evolved through scientific research in combination with trial and error. Handcrafted configuration has been needed to make deployments work in the real world. The difficulty in constructing a reliable large-scale system is that the complex pattern of interactions between components can produce behavior that is unexpected and difficult to predict or control."

On a similar note, 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."

We can't just feed data into an algorithm and hope for useful insights to pop out. As **Figure 15** shows, we must apply a rigorous, scientific method to AI/ML, which starts with clearly defined problems, moves on to a testable hypothesis, and then designs experiments to test that hypothesis.

Figure 15: The Scientific Method Applied to ML

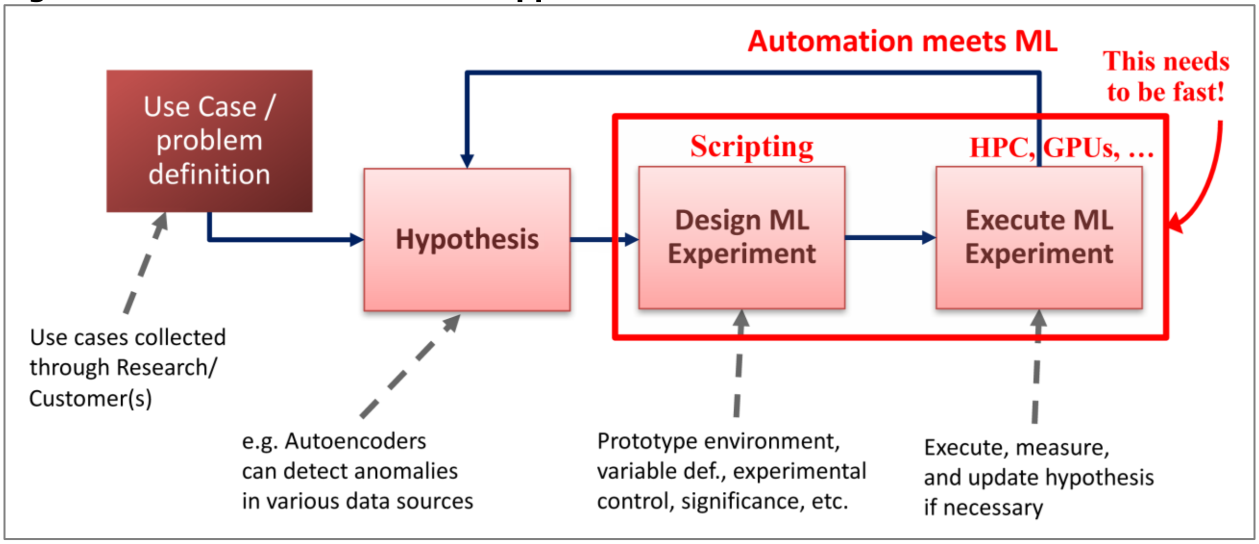

Source: David Meyer

6.4 Limitations of Tools

Only a small fraction of real-world ML systems is comprised of the ML code itself, as represented graphically by the small black box in the middle of **Figure 16**. The required surrounding infrastructure is vast and complex. As the paper Hidden Technical Debt in Machine Learning Systems reveals, Google finds it is common to incur massive ongoing maintenance costs in real-world ML systems.

Figure 16: ML Code Is Only a Small Part of the Job

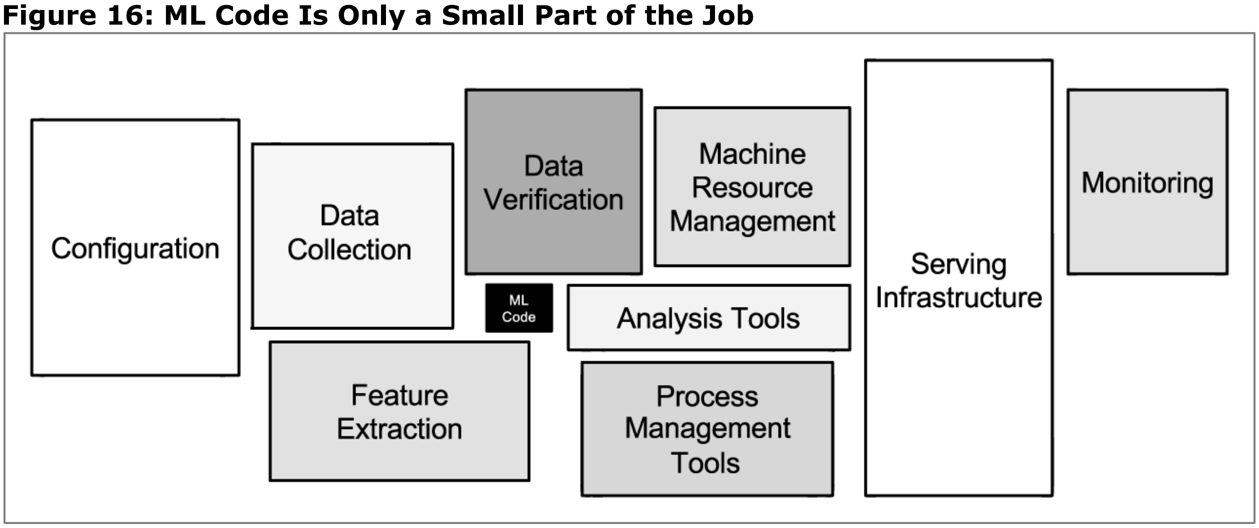

Source: Hidden Technical Debt in Machine Learning Systems, Sculley, et al.

To give a sense of the complexity of building a ML workflow for a networking application, consider **Figure 17**. First we must collect packet data, logs, KPIs, etc. using tools such as Wireshark, NFDUMP. Then the data needs to be ingested, cleaned, and normalized. The user then selects a model and trains, tests, and deploys the model. Only then does the user actually start to use ML algorithms on the data.

Figure 17: ML Workflow/Pipelines for Networking

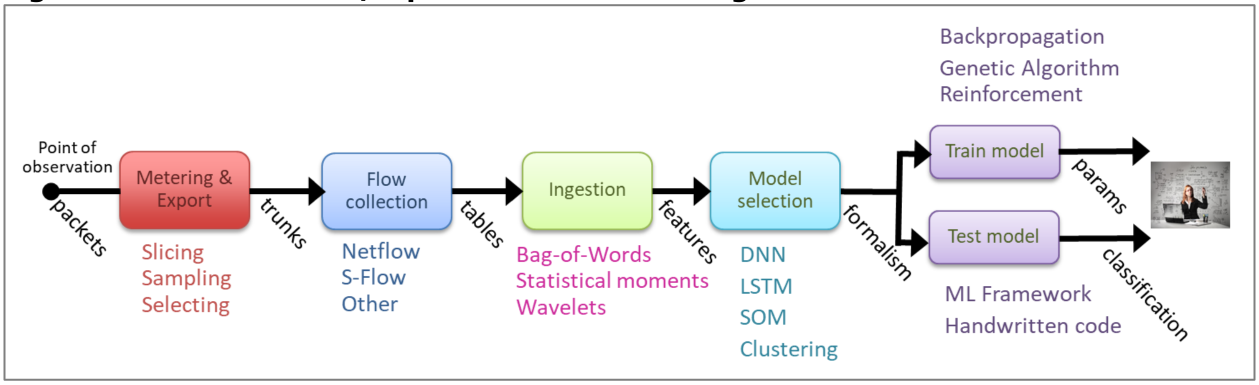

Source: David Meyer

A similar workflow is shown in **Figure 18**, courtesy of Deutsche Telekom.

Figure 18: ML in Deutsche Telekom's RT-NSM Architecture

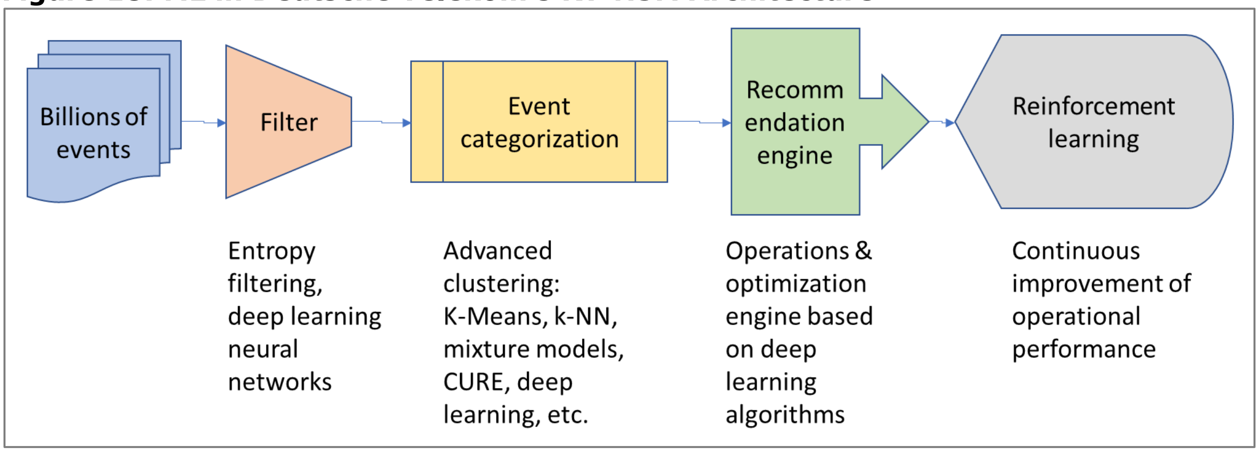

Source: Deutsche Telekom

7. ACADEMIC, SDO, CONSORTIA & OS INITIATIVES

Below we discuss several AI initiatives taking place in academia (Knowledge-Defined Networking), in standards organizations (ETSI and IEEE), industry consortia (TIP) and open source (Acumos).

7.1 Academia – Knowledge-Defined Networking

Knowledge-Defined Networking is a 2016 paper authored by academics from the Universities of Catalunya and Berkeley (California), as well as representatives from CA Technologies, Brocade, HPE, Intel, NTT and Cisco. In it, they explore the reasons for the lack of adoption of the Knowledge Plane proposed by D. Clark, a Senior Research Scientist at the MIT Computer Science and Artificial Intelligence Laboratory, and posit that the rise of two recent paradigms (SDN and network analytics) will facilitate the adoption of AI techniques in the context of network operation and control.

They describe a new paradigm that accommodates and exploits SDN, network analytics and AI, and provide use cases that illustrate its applicability and benefits. They refer to this new paradigm as Knowledge-Defined Networking.

Figure 19: Knowledge-Defined Networking

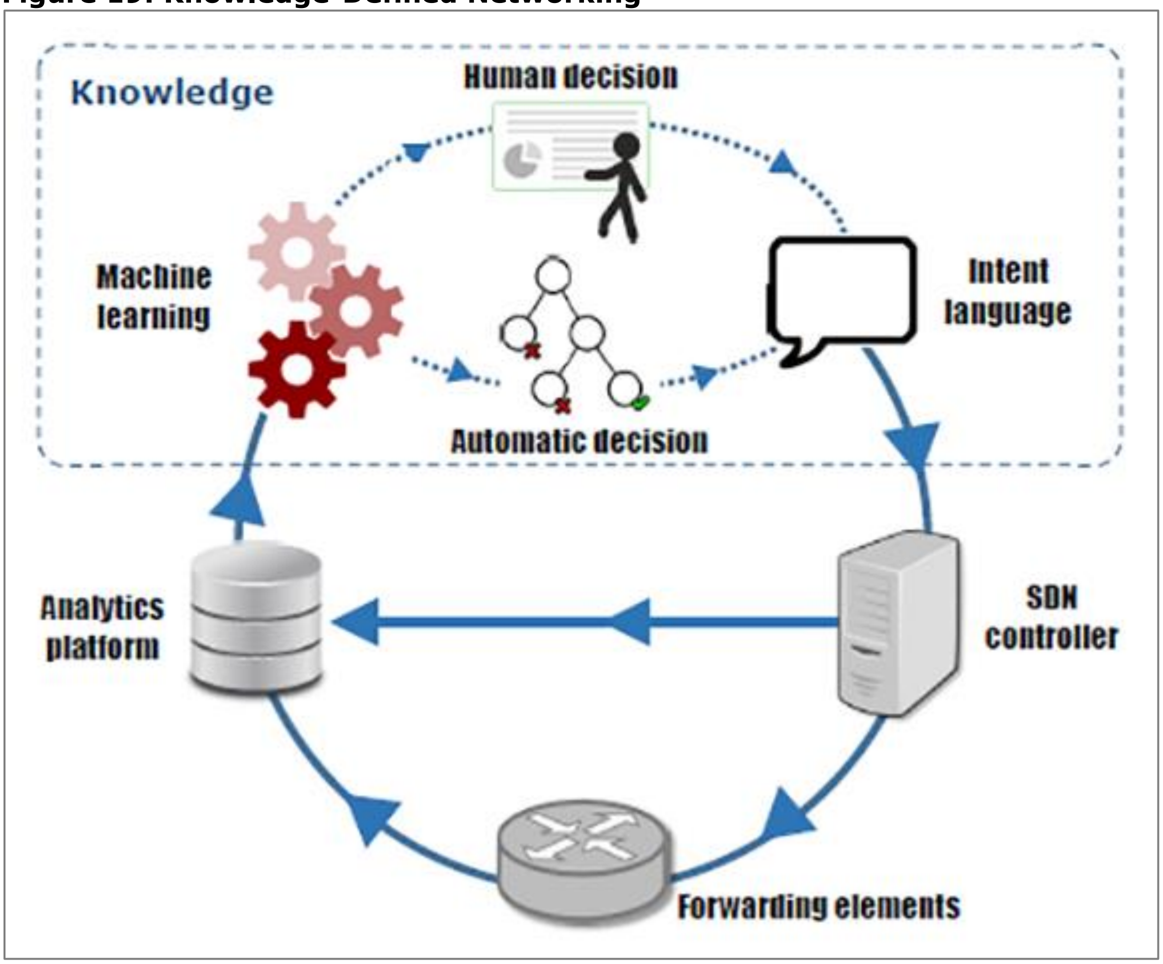

Source: Knowledge-Defined Networking, Mestres, et al.

The paper explains that one of the biggest challenges when applying ML to network operations and control is that networks are inherently distributed systems, where each node (switch, router, etc.) has only a partial view and control over the complete system. The emerging trend toward the centralization of control will ease this complexity. In particular, the SDN paradigm decouples control from the data plane and provides a logically centralized control plane, i.e., a logical single point in the network with knowledge of the whole.

In addition to the "softwarization" of the network, current network data plane elements, such as routers and switches, are equipped with improved computing and storage capabilities. This has enabled a new breed of network monitoring techniques, commonly referred to as network telemetry, which provide real-time packet and flow-granularity information, as well as configuration and network state monitoring data, to a centralized network analytics platform. In this context, telemetry and analytics technologies provide a richer view of the network compared to what was possible with conventional network management approaches.

7.2 Standards Development Organizations

ETSI's Experiential Networked Intelligence Group

ETSI believes that the use of AI techniques in the network supervisory system could help solve some of the problems of future network deployment and operation. It has therefore set up a new Industry Specification Group (ISG) on Experiential Networked Intelligence (ENI) to develop standards for a network supervisory assistant system.

The ENI ISG's initial white paper, Improved operator experience through Experiential Networked Intelligence (ENI), was published in October 2017. It included **Figure 20**, showing a progression from centralized control to AI-enabled, self-learning networks.

Figure 20: Future Evolution of Network Intelligence on Enhancing Experiences

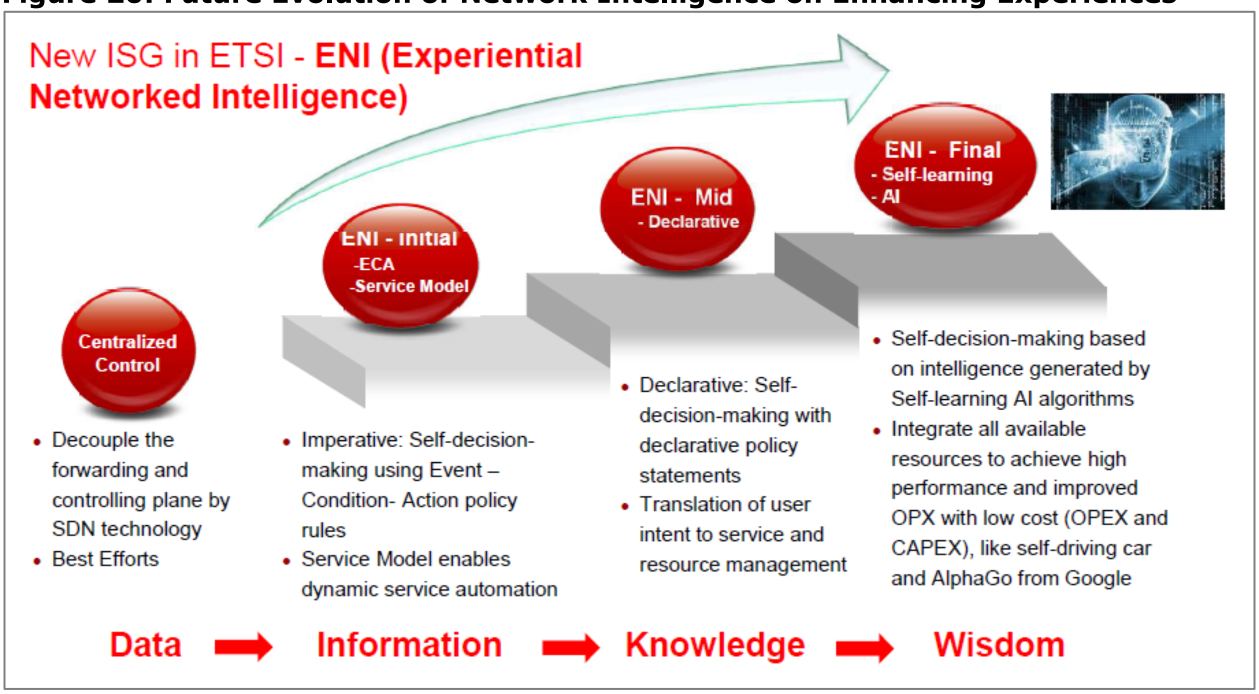

Source: ETSI

ENI followed its initial white paper with some very detailed use cases in the areas of infrastructure management (maintenance and planning), network operations (run-time), service orchestration and management (order to activation, SLA management, etc.), and assurance (monitoring and prediction). Examples include VoLTE optimization, network slicing for IoT security and SD-WAN management.

ENI has also published a requirements document covering service and network aspects (service orchestration and management, network planning and deployment, network optimization, resilience and reliability, security and privacy), functional aspects (data collection and analysis, policy management, data learning, interworking), and other aspects (performance, operational, regulatory, policy). This rather dry document will inform the system architecture specification that is due to be published next year.

Dr. John Strassner of Huawei Technologies has written the Context-Aware Policy Management Gap Analysis that draws on the work he is also doing for the MEF's Policy-Driven Orchestration (PDO) initiative, although "it is likely that ENI will need additional RPs [interface Reference Points]."

In terms of the gap analysis, the MEF's PDO fully meets all of ENI's requirements, except the use of AI to construct policies – however, "this is not hard to do because the MEF PDO information model uses software patterns," says Strassner. The IETF's Simplified Use of Policy Abstractions (SUPA) fully or partially meets most of ENI's requirements, whereas the TM Forum's SID satisfies hardly any.

The ENI ISG's Proof of Concepts Framework document provides a blueprint to which the PoCs must adhere. The first PoC will be championed by China Telecom, which was a major contributor to the original white paper – five of the 26 contributing authors were from the Chinese operator – as well as the new use case and requirements documents.

IEEE's Network Intelligence Emerging Technologies Initiative

The goal of the Institute of Electrical and Electronics Engineers' (IEEE) Network Intelligence Emerging Technologies Initiative is to research ways to embed AI in future networks to allow better agility, resiliency, faster customization and security.

The initiative argues that embedding intelligence into the network will provide a greater level of automation and adaptiveness, enabling faster deployment (from months down to minutes), dynamic provisioning adapted to the nature of network functions, and end-to-end orchestration for coherent deployment of IT and network infrastructures and service chains. The Initiative says it will also result in higher resiliency and better availability of future networks and services.

The group met in February 2018 and is participating in the call for papers for the IEEE's Journal on Selected Areas in Communications, which is expected to be published in the first quarter of 2019.

7.3 Industry Consortium – Telecom Infra Project

The Telecom Infra Project is a collaborative community of networking experts drawn from more than 500 member organizations (it is free to join) and with a board comprised of executives from BT, Deutsche Telekom, Facebook, Intel, Nokia, Telefónica and Vodafone.

In November 2017, TIP announced the launch of the Artificial Intelligence and Applied Machine Learning Project Group, co-chaired by Deutsche Telekom and Telefónica.

The objective of this project group is to define and share reusable, proven practices, models and technical requirements for applying AI and ML to reduce the cost of planning and operating telecom networks, and to understand and leverage customer behavior, optimizing service quality for an improved experience. The project group will collaborate across three work streams, as shown in **Figure 21**.

Figure 21: TIP ML Project Work Streams

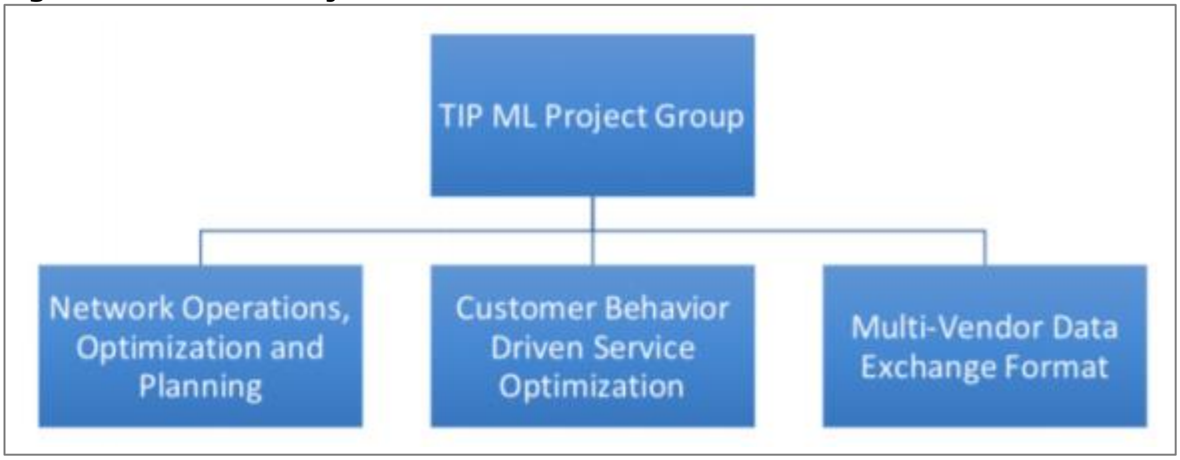

Source: TIP

1. Network operations, optimization and planning:
   - Predictive maintenance to identify faults before they impact network performance or customer experience
   - Automated recovery processes and avoidance of failures caused during provisioning and activation
   - Dynamic resource allocation and proactive maintenance via autonomous scheduling and configuration
2. Customer behavior-driven service optimization:
   - Predict customer behavior to help optimize network performance
   - Project best outcomes in bandwidth-intensive, latency-sensitive, and/or data-heavy applications
   - Analysis of needs by customer segment
3. Multi-vendor data exchange formats:
   - Select and adapt existing common data exchange formats that enable the network operations and customer behavior workstreams
   - Develop methods to minimize the handling and conversion of different vendor and network operator data formats
   - Unify data model definition and the meaning of individual data attributes within the multi-vendor data exchange formats

7.4 Open Source – Acumos

There are many open source projects related to AI, including Caffe, Keras, Scikit-learn, TensorFlow, Theano and Torch. The project that is most closely related to the telecom sector, in our view, is Acumos, which is the first project hosted by the Linux Foundation's Deep Learning group, the founding members of which are (in alphabetical order): Amdocs, AT&T, B.Yond, Baidu, Huawei, Intel, Nokia, Orange, Red Hat, Tech Mahindra, Tencent, Univa and ZTE.

Acumos is essentially an AT&T lab project that has been spun out to the Linux Foundation in order to apply the concept of crowdsourcing to software development. The beta version of Acumos has code from AT&T and Tech Mahindra and is available to download now.

The Acumos white paper explains how the development and deployment of AI applications is currently highly time-consuming and requires expensive, specialist talent. Acumos will provide a common framework that reduces the need for ML "rocket scientists" and accelerates development, thereby lowering the barriers to AI for CSPs and companies in other industries too.

With Acumos, general data scientists (who are not also ML specialists) will be able to create AI models built from tools such as Google's TensorFlow, Scikit-learn, R-Cloud or H2O (or directly programmed in Java, Python or R). Acumos will then package the model into a containerized microservice so that it can be easily integrated into applications by ordinary software developers without the need for a master's degree in data science and a specialization in AI development tools.

What makes Acumos even more interesting is that it borrows from the app store concept by providing a marketplace where academics and commercial data scientists can try to sell specialized ML models, customized to address a specific industrial use case. CSPs (and other companies) can then experiment with those crowdsourced models on their own datasets and, if they prove useful, pay to use them in production.

Acumos has taken the shrewd decision to leave data storage and processing outside its remit so that it is not exposed to any data privacy risk issues. CSPs will need to use their existing big data platforms (Kafka, Hadoop, Spark, map reduce, NoSQL, etc.) and feed their chosen models with this data within their own environment.

The specific areas of investigation discussed in the white paper include:

- Networking – responding to traffic anomalies in the real time (e.g., spinning up servers to meet demand spikes)
- Infrastructure – more accurate capacity planning
- Security – learning new attack signatures on the fly
- Customer care – (more) intelligent chatbots

The first release of Acumos (called Athena) is due in November, by which time hopefully there will be some more apps in the marketplace. Currently there are just four from AT&T (image classification, face detection, face privacy filtering and image mood classification) and two from Tech Mahindra (cross selling and customer segmentation). The current paucity of apps/models is partly because the monetization details (licensing) haven't yet been implemented and because there is not yet an automated mechanism for scanning models for malware.

Figure 22: Four Stages of AI Development Supported by Acumos

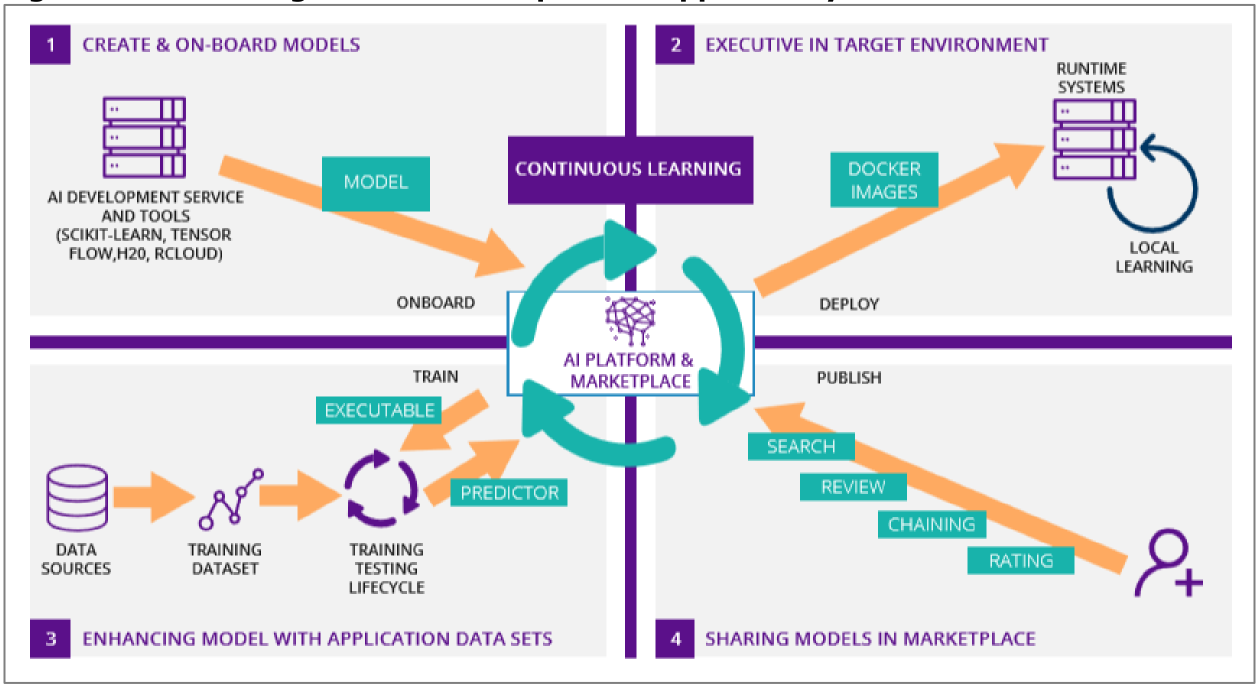

Source: Acumos

8. VENDOR PROFILES

Below are short profiles on some of the vendors we have come across in our research with AI-based offerings focused on the telecom industry. The use cases are typically in customer care, marketing and networking. Other use cases include IT operations, fraud and security.

Figure 23: Vendor Profile Summary

Company	Customer care	Marketing & CRM	Networking & IT Ops	Fraud & security
Afiniti	✓			
AIBrain	✓			
Anodot	✓		✓	✓
Arago			✓	
Aria Networks			✓	
Avaamo	✓			
B.Yond			✓	
Cardinality		✓	✓	
Guavus	✓	✓	✓	✓
Intent HQ		✓		
IPsoft	✓			
Nuance	✓			
Skymind			✓	✓
Subtonomy	✓		✓	
Tupl	✓		✓	
Wise Athena		✓		✓

8.1 Afiniti

Founded in 2006 and based in Washington, D.C., Afiniti employs approximately 650 people. Afiniti Enterprise Behavioral Pairing uses AI to identify subtle and valuable patterns of human interaction to pair individuals based on behavior, leading to more successful interactions and measurable increases in enterprise profitability. Afiniti's customers in the telecom and entertainment markets include T-Mobile USA, Bouygues Telecom, Virgin Media and Sky UK.

Afiniti initially deployed its technology within both of T-Mobile USA's outsourced call centers, replacing their historical "first-in first-out" call flow model. Sales conversion rates across T-Mobile's 1,000 tele-sales representatives jumped by more than 5 percent. T-Mobile did not bear any upfront investment. Afiniti accepted the entire commercial risk in exchange for a commissioned engagement in which it received a fixed amount for each incremental customer delivered.

After nine months with Afiniti, T-Mobile elected to deploy Afiniti's technology across its entire contact center estate, comprising 10 PBX/ACDs and 20,000 agents. Afiniti claims that it generates more than $150 million a year in incremental cash flow for T-Mobile.

8.2 AIBrain

Founded in 2012 and based in Menlo Park, Calif., AIBrain employs approximately 25 people. AIBrain combines several technologies to deliver AI to customer service centers through an SDK/API:

- AICoRE is a cognitive reasoning engine, simulating human intelligence across the full spectrum of problem solving.
- GCA achieves human-like conversation.
- Memory Graph is a human-like AI memory system, seamlessly integrating episodic and semantic memories for all types of intelligent agents.

8.3 Anodot

Founded in 2014 and based in Ra'anana, Israel, Anodot employs approximately 80 people. Anodot is a real-time analytics and automated anomaly detection system that discovers outliers in time series data and turns them into valuable business insights. Using ML algorithms, Anodot isolates issues and correlates them across multiple parameters in real time, supporting business decisions through its uncovered insights. Telecom customers include Magyar Telekom (a Deutsche Telekom subsidiary) and Optus.

In the telecom vertical, Anodot provides solutions for the following areas:

- Service quality – automatically tracks and correlates KPIs from multiple sources such as customer behavior and network data.
- Roaming – identifies issues with subscriber event processing, account details, and billing by detecting and diagnosing anomalies.
- Revenue assurance – identifies revenue leakage and surfaces discrepancies between expected results and how events are actually billed.
- Customer care – extracts real time insights from customer usage patterns by collecting and analyzing data from billing systems, network operations, subscriber accounts and customer care systems.

8.4 Arago

Founded in 1995 and based in Frankfurt, Germany, Arago employs approximately 120 people. Arago was originally founded by algorithm expert Chris Boos as a research company. Since 2014, private equity firm KKR has supported Arago's international expansion. Arago offers a general AI platform, HIRO, that can be used for IT and business process automation.

HIRO can be implemented across general business functions such as finance and HR. The company also provides a solution for IT operations that it claims can automate 87 percent of the IT stack, decreasing operational cost by 50 percent. The IT operations solution is comprised of the following modules:

- **Compute Infrastructure Operations:** autonomously runs IT infrastructure management processes within the data center's compute infrastructure, networks, middleware, platform, and application layer. Use cases could be: incident management, proactive triage, root cause analysis, change and configuration management, capacity management, performance management.

- **Service Desk Operations:** autonomously runs service request management processes across the entire enterprise IT landscape. The solution proactively realizes the authentication; validation; and execution of user service requests, such as: password reset; user account management; mailbox provisioning; application access administration; software provisioning.
- **Security and Compliance Operations:** security incident and event management; management of security alerts; automatic remediation of security issues in near-real time; security policy enforcement and security audit trail; root cause analytics.
- **Application Support:** automates maintenance and administration processes across the entire application lifecycle: application health check; application performance diagnostics; incident remediation; database administration; security compliance audit; software provisioning; SSL certification provisioning.

Figure 24: Arago's HIRO Platform

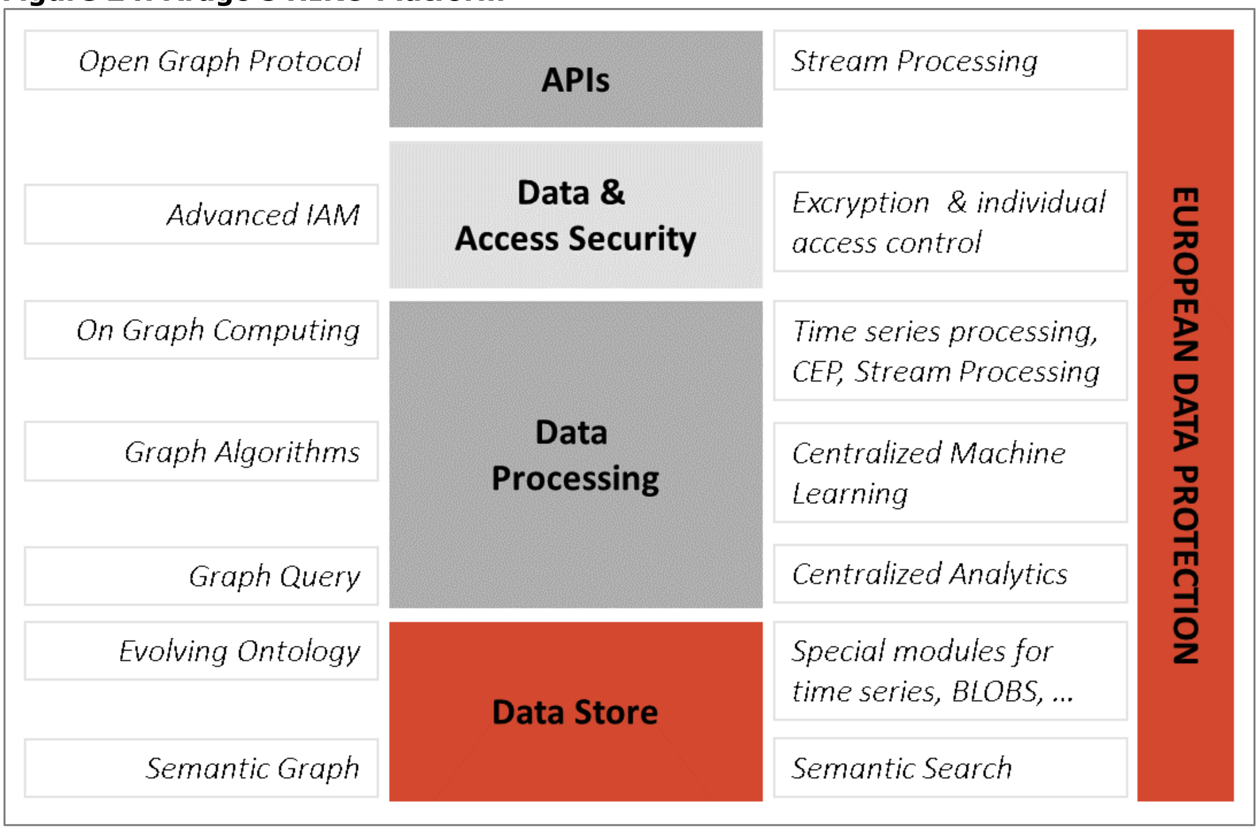

Source: Arago

8.5 Aria Networks

Founded in 2005 and based in Bath, U.K., Aria Networks employs approximately 30 people. Aria Networks offers a range of AI-powered solutions addressing:

- Network visualization
- Traffic modeling and prediction; routing analysis; failure impact assessments
- Vulnerability analysis and network optimization

- What-if analysis
- Margin and contract analysis

Aria Networks' solutions can be applied to SDN/NFV rollout, SD-WAN optimization, network planning and data center automation.

Figure 25: Aria Networks AI Engine Within Next-Generation OSS

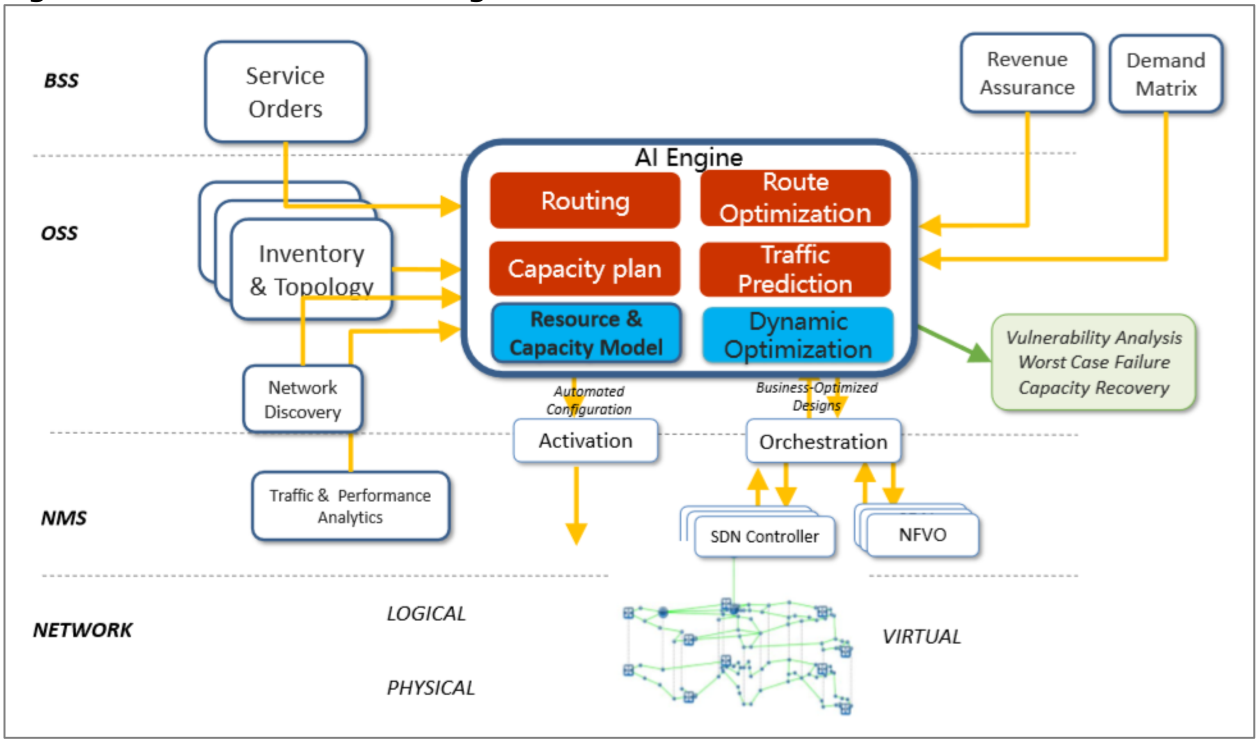

Source: Aria Networks

8.6 Avaamo

Founded in 2015 and based in Los Altos, Calif., Avaamo employs approximately 55 people. Avaamo makes software that specializes in conversational interfaces using neural networks, speech synthesis and deep learning to make conversational computing for the enterprise.

Avaamo's CSP customers are using conversational AI to guide their subscribers through the bewildering choice of phone models, prepaid and postpaid plans, and add-on offers, while maximizing the value for their subscribers and driving higher ARPU and profit margins.

Avaamo CSP use cases include:

- **Provisioning/billing:** activating plans, understanding data usage, doing top-ups, and paying off bills for individual subscribers,
- **NOC support:** conversational applications to understand network utilization, latency, and error rates for wholesale and business customers,
- **Field service:** Proving field service personnel with latest information of preventive maintenance and service ticket completion.

- **Call center agent support:** conversational applications on the agent desktop to assist in understanding customer questions.
- **Contact center/customer support:** When a CSP has to support more than 2,000 models of phones, modems, routers, etc., the technical skills required of the customer support staff go up as do the average handle times. Avaamo's customers are using conversational AI to deflect calls from voice to web and mobile chat channels, thereby improving average handle time, first contact resolution, and customer satisfaction/NPS scores.

8.7 B.Yond

Founded in 2017 and based in Cupertino, Calif., B.Yond employs approximately 130 people. B.Yond is an AI company that develops automation software for CSPs seeking to capitalize on the move to 5G and Edge Cloud. B.Yond's INFINITY platform leverages AI, combining real-time, predictive analytics with automation and enabling self-provisioning, self-optimizing and self-healing networks. B.Yond currently has two products based on the INFINITY platform:

- INTEGRITY – Asset & Inventory Reconciliation. Provides precise and reliable views of network data for transforming your network or implementing operational efficiencies.
- SANITY – Smart Alarm Management. Offers AI functionality for dealing with exponential increases in network complexity.

Figure 26: B.Yond INFINITY Platform High-Level Architecture

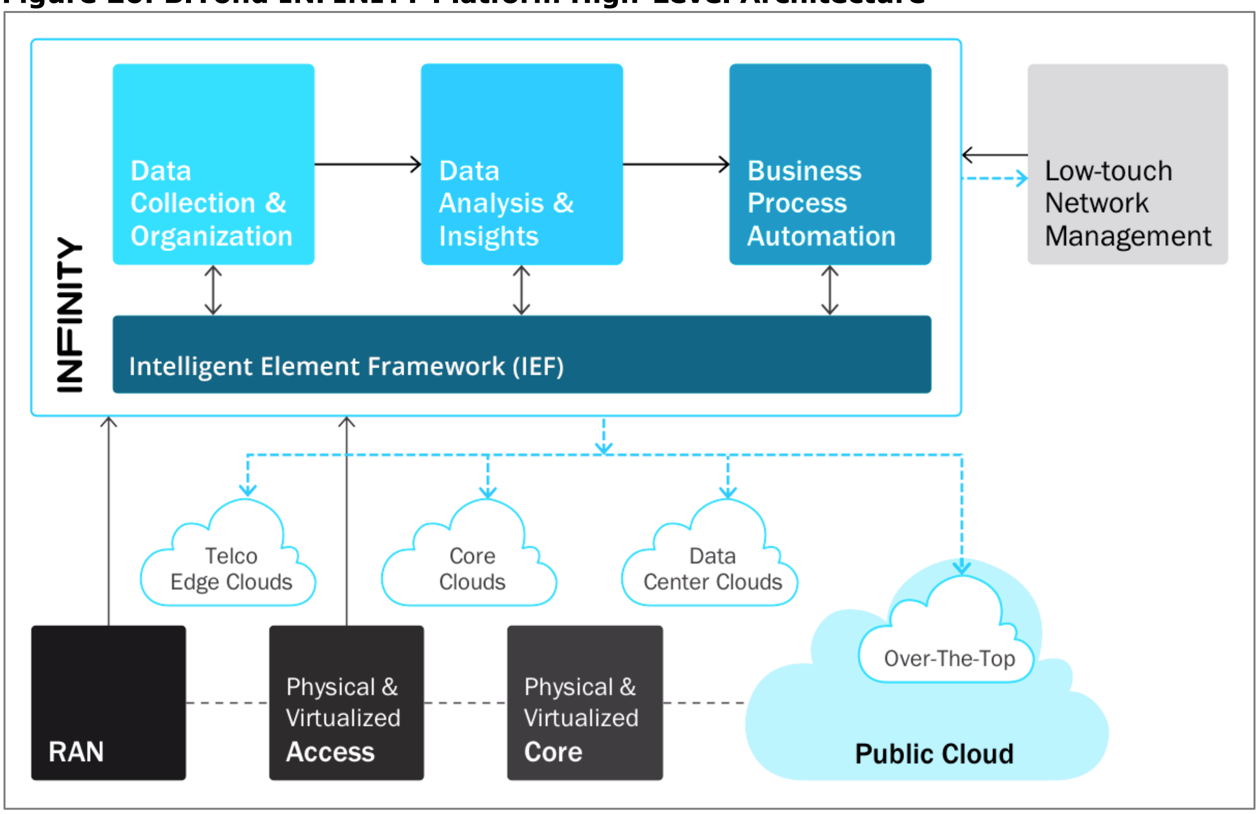

Source: B.Yond

8.8 Cardinality

Founded in 2015 and based in Guildford, U.K., [Cardinality] employs approximately 20 people. Cardinality offers analytics tools that can be applied to:

- Radio cell management
- CEM
- Operational intelligence
- Device analytics
- Marketing insights

Perception is Cardinality's full stack offering, which incorporates ETL Engine and Analytics Engine in a single coordinated platform. It offers a focused, generic solution across the telecom industry. Perception is proven to scale – from large enterprise deployments down to initial PoCs, which can cover all aspects of scale.

The platform concentrates on extracting the maximum value from our customers data. It handles every task from the initial ingestion process from any data source, cleaning and preparing diverse data, post ingestion enrichment and KPI creation, through to serving data out for operational use and analysis.

Perception doesn't penalize you financially for data growth, as licensing is not based on data volume. It also offers a better return on investment than standalone Hadoop deployments.

Figure 27: Cardinality Perception Analytics System

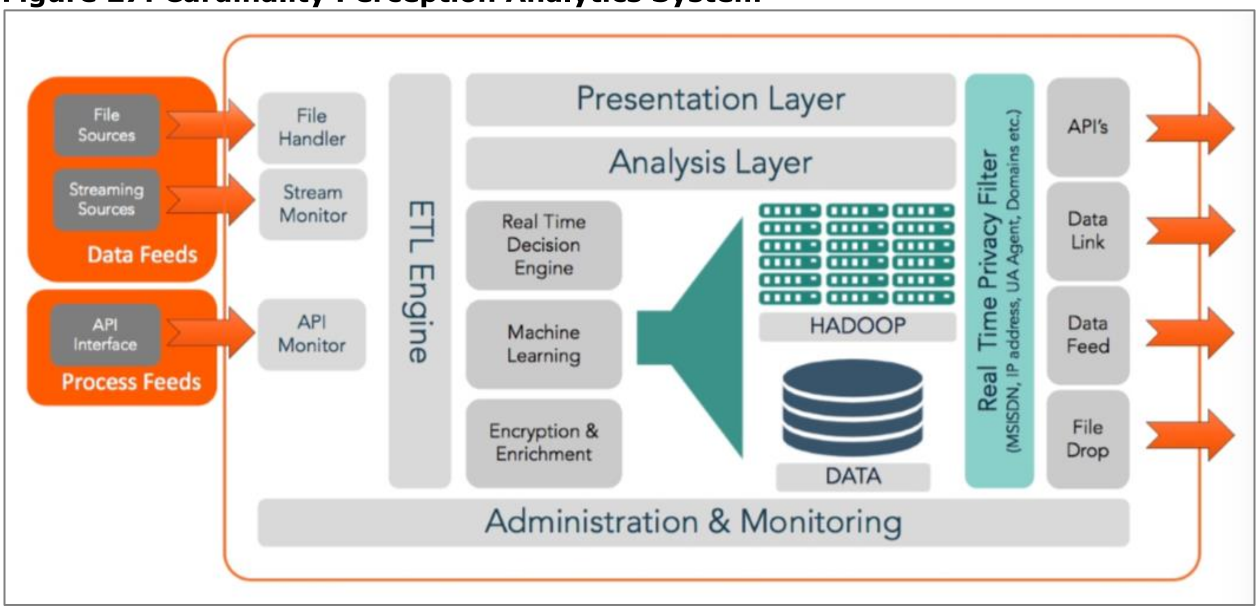

Source: Cardinality

8.9 Guavus

Founded in 2006 and based in San Jose, Calif., [Guavus] employs approximately 250 people. Guavus was acquired by French aerospace, transport, defense and security vendor Thales in 2017 for up to $215 million in cash. Thales expected Guavus revenues in 2017 to exceed $30 million.

Guavus's products are based around its Reflex platform, which comprises a big data Base Processing Layer and the Reflex Analytics Fabric. The Reflex Analytics Fabric enables the development of big data and AI/ML analytics applications through its Analytics Engine, App Accelerators and Micro Apps.

Customers can leverage the Reflex Analytics Fabric to build their own analytics apps without needing to be big data or analytics experts. This layer includes an extensible set of ML algorithms to which Guavus has added AI or "reasoning." AI enables the Analytics Engine to draw logical inferences even when presented with patterns of events that have not been seen before.

Figure 28: Guavus Reflex Platform

Source: Guavus

Guavus's products include:

- **Live Ops:** adaptive analytics to correlate separate data sources (call logs, trouble tickets, service calls, network and equipment change records, etc.) to identify and understand customer impacting events in real time and then recommend steps to repair.
- **Proactive Ops:** proactively anticipate events that may cause network problems, identify which ones will have the biggest customer impact and take automated actions as needed.
- **Security Intelligence:** automatically detects anomalous behavior to show security analysts where threats may be imminent, without overwhelming them with false positives.

- **Smart Care:** integrates with existing customer care systems and recommends resolutions through advanced predictive algorithms and AI.
- **Marketing Insight:** dynamic customer segments can be created based on particular campaigns and targeted offers presented in real time to increase acceptance rates.
- **Smart Industry & IoT:** an out-of-the-box solution that can run at the edge or in a central location to automatically pinpoint customer behaviors and preferences, classify asset usage, identify performance issues and root causes, and take closed-loop actions.
- **Alarm IQ:** harnesses the power of AI to eliminate alarm noise without requiring changes to NOC operator workflows.

Guavus supplies six of the seven top telecommunications providers and three of the four top MSOs globally.

8.10 Intent HQ

Founded in 2010 and based in London, Intent HQ employs approximately 36 people. Intent HQ takes customer data and organizes it into actions (something you do), attributes (something you are) or entities (something you use). It then transforms the raw data into human-like profiles that represent customers as people – their habits, interests and motivations. Using bespoke algorithms and its Intent Graph, it finds patterns in the data to see how customer behavior impacts business KPIs. Users can explore the results asking their own business questions, such as identifying which competitors are capturing the attention of Netflix-obsessed new parents, or understanding exactly who to target for the next iPhone launch.

Nina Bibby, CMO for Telefónica UK, says "Our partnership with Intent HQ has helped us know our customers on a completely different level, as people. This is transformational, allowing us to create the relationship of equals we strive to have with customers and strengthening our market position. For us, data is at the heart of every decision we make."

According to the article Telco Data Analytics Europe: Key Takeaways, Intent HQ CEO Jonathan Lakin believes AI can help CSPs build human-like customer profiles that will anticipate what customers want to do, interpret their intent and automate actions to deliver a personalized experience – a customer data hub/cloud platform that takes all of customer interactions, interests, habits and behaviors, based on visits to websites, etc., ingesting at event level. The giant data set can then be tuned with AI algorithms to micro-segment customers, reduce the cost to serve certain segments, run data-driven marketing campaigns, etc.

Lakin believes that most CSPs are still figuring out how to move to being data-driven, and that the current ways of working in silos must dramatically change. One of the first things his company does with any operator is spend time with each department to determine maturity in terms of using data intelligently to improve the business processes, interactions and data sharing between groups to drive common business goals.

8.11 IPsoft

Founded in 1998 and based in New York, IPsoft employs approximately 2,500 people. IPsoft provides AI solutions for enterprises (e.g., banking, insurance, healthcare, telecom). Their flagship AI product, Amelia, is a fully conversational AI used for intelligent chat-based (text, not voice) interactions.

Amelia stores facts, concepts, and the associations between them in her semantic memory. From standard operating procedures to policy documents, she reads information and then applies it to conversations. She remembers every interaction she has and can use that information to deliver faster and more informed results. Amelia can dynamically navigate business process flows without having to follow a step-by-step process to achieve a desired outcome. This allows her to jump from one process to another, if a conversation requires her to do so. Amelia uses the Affective Computing and Sentiment Analysis techniques to computationally model user's emotion, mood and personality.

According to this article on IPsoft's website, Karine Brunet, Director of Technology Shared Services at Vodafone, said her company began using Amelia two years ago to support employees with IT issues. Vodafone is now using Amelia in seven markets in English, German and Spanish, and 58 percent of employees who've contacted the IT service desk have used Amelia. Of the 20,000 chats per month that Amelia handles, she's able to complete 53 percent without human intervention. Brunet notes, "We believe that in the next few months we're going to increase that level of autonomy, and we think we should be close to 65 to 70 percent in terms of autonomy."

Vodafone has integrated Amelia with 22 of the company's back-end information systems, including its Active Directory and ITSM ticketing system, and more recently SAP SuccessFactors, the company's HR system.

8.12 Nuance Communications

Founded in 1994 and based in Burlington, Mass., Nuance employs approximately 11,600 people. Nuance provides speech and imaging applications based on automatic speech recognition (ASR), natural language understanding (NLU), and text-to-speech (TTS). In fiscal year ending September 2017 Nuance's revenue was approximately $1.9 billion. In the telecom sector customers include French incumbent Orange and Vodafone Australia.

In the CSP market, Nuance offers:

- Contact Care Center Solutions – Nuance Mobile Care: Service From Mobile Handsets; IVR Self-Service Applications; Proactive Outbound Notification Solutions; CTI and Agent Desktop Solutions; Voice Verification Solutions
- Multimodal Messaging, Dictation and Search Solutions
- Voice to Text Services – for voicemail, missed call services and a variety of peer-to-peer messaging categories.

8.13 Skymind

Founded in 2014 and based in San Francisco, Skymind employs approximately 30 people. Skymind helps companies build deep-learning applications for media, images and sound, and time series data for finance, healthcare, telecom and IoT. CSP customers include China Unicom, KDDI, Orange and Vodafone.

Skymind's Intelligence Layer (SKIL) helps bridge the gap between the Python ecosystem and the JVM ecosystem with a cross-team platform for Data Scientists, Data Engineers and DevOps/IT. SKIL enables interoperability between data science and big data frameworks by standardizing and orchestrating AI workflows within a single, consolidated platform.

One CSP use case for Skymind is SIM box fraud. Orange Silicon Valley is working with Skymind to prevent SIM box fraud on its mobile network. Using an artificial neural network (ANN) called an autoencoder, Orange Silicon Valley analyzes call detail records (CDRs) to find patterns that identify fraud. The ANN predicts the likelihood that an instance is fraudulent. Where a static rule system flags cases as likely fraud or not based on preconceived decision trees, the ANN enables Orange's analysts to prioritize high probability cases of SIM box fraud adaptively.

Figure 29: Skymind's SKIL Platform

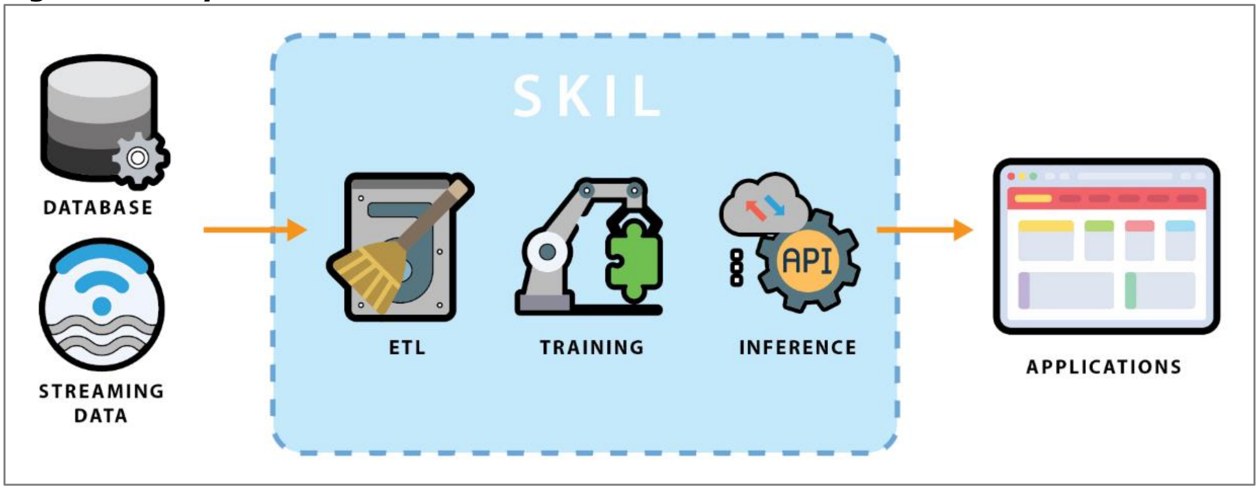

Source: Skymind

8.14 Subtonomy

Founded in 2012 and based in Stockholm, Subtonomy employs approximately 13 people. Subtonomy's platform gathers data in real time from different sources in mobile network and uses AI to detect performance deviations. The platform collects around 1,500 data points every day for each subscriber. Data sources can include passive probes, cell site data, CRM systems, user devices and even billing.

The platform uses AI to learn over time how the mobile network and its subscribers behave. This is used to automatically set alarm thresholds, know when to trigger certain events, and predict how subscribers will behave. Customers include Hutchison 3 (Sweden and Denmark), Telenor Sweden and Telia (Norway and Estonia).

Products include:

- SubSearch: enables MNO customer care teams to search for a single subscriber in real time, see service performance and potential issues, and get recommendations to solve issues.
- CorpDash: an enterprise SLA performance tracking tool for MNOs that filters actual traffic data by specific enterprise customer. CorpDash can predict which users are likely to be affected by certain issues, e.g., planned maintenance work.
- NocMap: used by network and operations teams, it gives a very detailed view of the network and aggregated views. NocMap is self-configuring on setup and the application will then learn the network behavior in order to detect abnormalities.

- Roamers: monitors all roaming traffic (inbound and outbound) to assure service quality and optimize offerings. Roamers is used by wholesale, customer care and operations.

Figure 30: Subtonomy Platform Overview

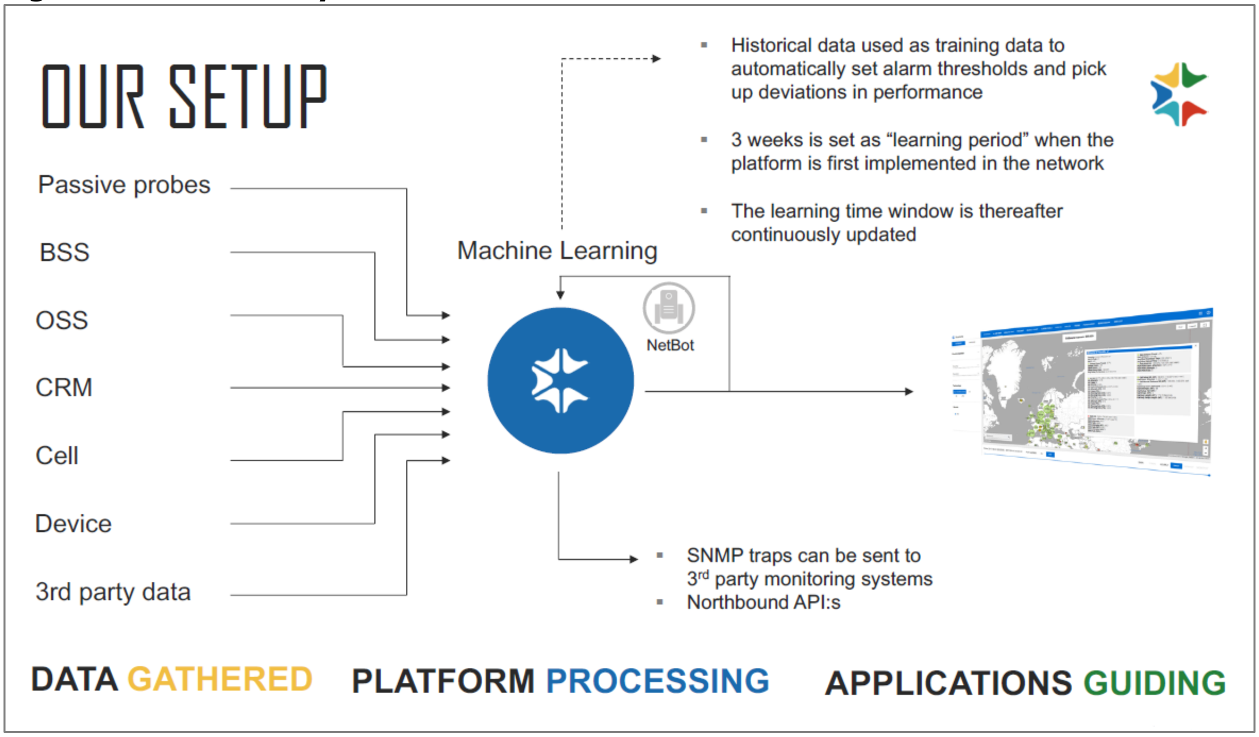

Source: Subtonomy

8.15 Tupl

Founded in 2014 and based in Bellevue, Wash., Tupl employs approximately 50 people. Tupl's AI engine, TuplOS, is a backend platform that facilitates the digitalization of engineering knowledge to promotes process automation.

TuplOS is positioned as an augmentation for expensive, manually operated OSS systems. TuplOS addresses the following areas:

- Automatic customer complaint resolution – integrating with network and subscriber performance data, the platform can reduce network engineers' workload and help support teams quickly identify and resolve end-customer issues
- Network advisor – an AI-powered assistant designed to boost network performance team productivity by automating daily repetitive performance and troubleshooting tasks, and providing orchestration between existing network optimization or planning tools
- Tupl Unifier – collects, manages and visualizes performance, call records and internal databases

CSP customers include T-Mobile USA and Softbank Japan.

Figure 31: Tupl in Context

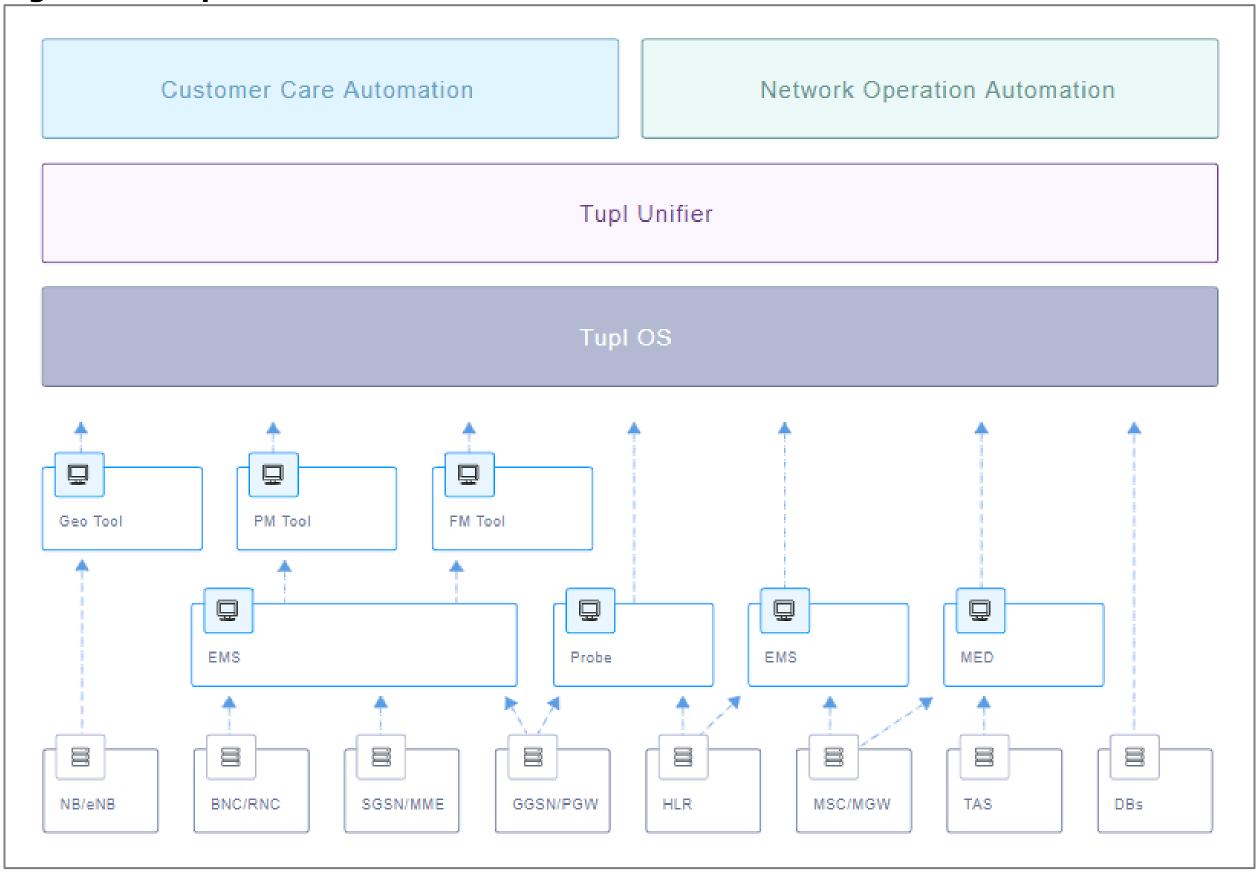

Source: Tupl

8.16 Wise Athena

Founded in 2013 and based in San Francisco, Wise Athena employs approximately 20 people. Wise Athena provides a suite of pricing and trade promotion optimization software: constraint-based predictive optimization, optimization through iterative scenarios, prediction of "best" price and promotion option, cannibalization and price ratio effects.

The platform leverages AI to help consumer companies predict the right price in just one click. Wise Athena analyzes thousands of data points, including competitor actions, to choose the best pricing, thereby maximizing sales and profit during every promotional cycle.

Wise Athena has also applied AI to telecom operator churn prediction and CSP fraud.

9. CONCLUSIONS

AI and ML are not new topics. Even in the context of telecom, as far back as 1993 researchers were exploring AI techniques that they thought would be essential to the transformation of the telecom network. The key factors that have led to an acceleration in progress in ML in recent years include breakthroughs in neural network theory, the availability of massive data sets for academics to experiment with, and the rise of public cloud (AWS, Azure, etc.) making computing capacity readily available and cheap.

Key AI use cases in telecom include:

- **Network operations monitoring and management:** AI and ML approaches are beginning to emerge in the networking domain to address the challenges of virtualization and cloud computing. Increased complexity in networking and networked applications is driving the need for increased network automation and agility.
- **Predictive maintenance:** A Heavy Reading 2017 survey of CSPs found that predictive maintenance was the top use case for ML in telecom.
- **Fraud mitigation:** Fraud detection and prevention was the fifth most popular use case in the same Heavy Reading survey. According to the Communications Fraud Control Association, fraud costs the global telecom industry $38 billion annually.
- **Cybersecurity:** Security was the second most popular use case in our survey. Heavy Reading's Telecom Security Market Tracker has found that there is guarded optimism over AI for the automation of CSP security.
- **Customer service and marketing virtual digital assistants:** One of the key applications of AI/ML in the telecom sector to date has been the use of chatbots to augment or replace human call center agents. Seven of the 10 CSP profiles in this report include a discussion of their use of AI in customer care.
- **Intelligent CRM:** AI can be applied to CRM in areas such as personalized promotions, cross-sell/up-sell opportunity identification, and churn prediction and mitigation.
- **CEM:** As digital touchpoints continue to grow, analytics and AI are essential tools for CSPs to understand the health of the network, the customer journey (customer care, billing, etc.) and real-time service quality.

Of the 10 CSP case studies of AI profiled in this report, eight of them describe using AI in networking, seven in customer care and two in fraud/security. Of the 16 vendors profiled in this report, nine of them are applying AI to networking, nine to customer care, four to marketing/CRM, and four to fraud/security.

A 2017 survey by TM Forum found that 52 percent of respondents claimed that they were already using ML and analytics for network management, and a further 38 percent planned to do so in the next two years. In a Heavy Reading 2017 survey, most respondents said that AI/ML would become a critical part of network operations by 2020.

The greatest challenges to applying AI/ML to networking (as opposed to more generic applications within CSPs, such as CRM or chatbots) are:

- Data that is dirty, unavailable or difficult to access
- Lack of data science talent

- Lack of a clear question to answer
- Limitations of tools

These challenges may be addressed by initiatives currently ongoing in the industry, including:

- ETSI's Experiential Networked Intelligence Group
- TIP's AI and Applied ML Group
- The Linux Foundation's Acumos open source project

Although Heavy Reading is optimistic about the potential for AI/ML in telecom, we are also aware that there is a great deal of hype around the terms with many industry participants using these acronyms inappropriately in their marketing when, in fact, their systems are still based on a traditional rules-based approach.

Nonetheless, an increased level of automation is required to manage virtualized networks (including 5G), and AI/ML could play an important role, particularly in supporting real-time decision-making. The complexity of communications networks seems to rise inexorably with the deployment of new services, such as SD-WAN, and new technologies, such as SDN/NFV.

To meet ever-rising customer expectations, CSPs must increase the intelligence of their network operations, planning and optimization. ML and AI will be key to automating network operations and optimizing the customer experience. ML and AI promise to reveal new insights from network telemetry and flow data, enabling CSPs to predict capacity demands and scale their networks appropriately.

Heavy Reading envisages the adoption of ML falling into three phases: learning, advising and autonomous.

- In the learning phase, the system observes events and incidents and notes the actions taken by operations staff to address them.
- In the advising phase, the system recognizes learned events and incidents, and makes recommendations on how to resolve the problem. If operations staff do not follow these recommendations, this is useful feedback to the system that can enable it to improve its future recommendations.
- In the autonomous phase, the system can apply its recommended changes automatically without human intervention. This clearly requires some connection between the ML system and the network management and orchestration systems. The autonomy will be confined to known event signatures that occur often and have well-understood mitigation processes. Other events will fall back to the advising or learning phase.

Care must be taken with the introduction of ML systems that operations staff do not perceive these new technologies as a threat to their employment. AI/ML should be seen as a means to free up the time of staff to focus on more complex tasks, rather than responding to alarms. 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. The management of telecom networks is a multifaceted process that current ML technology can only partially address.

TERMS OF USE

LICENSE AGREEMENT
This report and the information therein are the property of or licensed to Heavy Reading, and permission to use the same is granted to purchasers under the terms of this License Agreement ("Agreement"), which may be amended from time to time without notice. The purchaser acknowledges that it is bound by the terms and conditions of this Agreement and any amendments thereto.

OWNERSHIP RIGHTS
All Reports are owned by Heavy Reading and protected by United States Copyright and international copyright/intellectual property laws under applicable treaties and/or conventions. The purchaser agrees not to export this report into a country that does not have copyright/intellectual property laws that will protect Heavy Reading's rights therein.

GRANT OF LICENSE RIGHTS
Heavy Reading hereby grants the purchaser a non-exclusive, non-refundable, non-transferable license to use the report for research purposes only pursuant to the terms and conditions of this Agreement. Heavy Reading retains exclusive and sole ownership of all reports disseminated under this Agreement. The purchaser agrees not to permit any unauthorized use, reproduction, distribution, publication or electronic transmission of this report or the information/forecasts therein without the express written permission of Heavy Reading.

DISCLAIMER OF WARRANTY AND LIABILITY
Heavy Reading has used its best efforts in collecting and preparing this report. Heavy Reading, its employees, affiliates, agents and licensors do not warrant the accuracy, completeness, currentness, noninfringement, merchantability or fitness for a particular purpose of any material covered by this Agreement. Heavy Reading, its employees, affiliates, agents or licensors shall not be liable to the purchaser or any third party for losses or injury caused in whole or part by Heavy Reading's negligence or by contingencies beyond Heavy Reading's control in compiling, preparing or disseminating this report, or for any decision made or action taken by the purchaser or any third party in reliance on such information, or for any consequential, special, indirect or similar damages (including lost profits), even if Heavy Reading was advised of the possibility of the same. The purchaser agrees that the liability of Heavy Reading, its employees, affiliates, agents and licensors, if any, arising out of any kind of legal claim (whether in contract, tort or otherwise) in connection with its goods/services under this Agreement shall not exceed the amount the purchaser paid to Heavy Reading for use of this report.

DISPUTE RESOLUTION
This License will be governed by the laws of the State of New York. In case of a dispute arising under or related to this License, the parties agree to binding arbitration before a single arbitrator in the New York City office of the American Arbitration Association. The prevailing party will be entitled to recover its reasonable attorney fees and costs.

Heavy Reading
P.O. Box 1953
New York, NY 10156
Phone: +1 212-600-3000
www.heavyreading.com