Implementing Operational AI in Telecom Environments

Pablo Tapia, Enrique Palacios, Laurent Noël, Petri Hautakangas - Tupl Inc.

Abstract

The last decade has witnessed a tremendous transformation of the telecom market, placing operators into a complex situation: on one side, increased competition drives ARPU / ARPA down, while the race to deliver leading edge user experience and the ever increasing demand for data pushes operators to increase CAPEX. To sustain their financial growth, operators must improve their operations efficiency, and automation is their only choice. To maintain their competitive advantage, they must deploy innovative solutions that deliver superior user experience at minimal cost, with very fast development cycles. Machine Learning (ML) tools are key technology enablers that can help achieve these ambitious targets and challenges. But Artificial Intelligence (AI) is not magic, and applying AI to meet telecom-grade quality takes a lot more work and skills than making shopping recommendations based on a person's navigation history. This paper presents the key considerations required to deploy AI on an operator's infrastructure, and provides a real life example of how these promises can be met. The following case is believed to be the first published example of automation of complex telecom engineering decisions through the use of an AI platform and digitalization of engineering knowledge. The result achieves 90% process automation with 98% reduction in processing time, all while reducing the number of unresolved issues by a factor 4x. This demonstrates the potential of AI technologies, when combined with the proper domain-level expertise. Furthermore, a comprehensive knowledge digitalization strategy based on AI/ML has the potential to transform OPEX into CAPEX costs, and significantly improve the operator's bottom line.

Challenging Times for the Industry

The Telecom industry has lived through a tremendous technology evolution over the last couple of decades, fueled by the ever-increasing demand on data services both in mobile and in fixed networks. While we have witnessed quite impressive growth on the customer acquisition and network footprint expansion, the developed countries are now reaching a saturation point where telecom operator's prospects for revenue growth are under pressure due to increased competition. This competition is visible on multiple commercial fronts: price wars, unlimited data caps, time-limited offers of all sorts, bundle package discounts, etc. Inevitably, this competition also leads to an erosion of their ARPU.

In contrast to slowing subscriber growth and stagnant ARPU, traffic levels are expected to continue growing rapidly. At the same time, customers are growing more demanding in their quality expectations, and operators are also trying to differentiate themselves by providing a better user experience, which involves a combination of both content and quality of delivery. To summarize, it is clear that operators' strategic priorities are focused on managing the customer experience, and in a cost-effective manner. This was also evidenced in the survey results below.

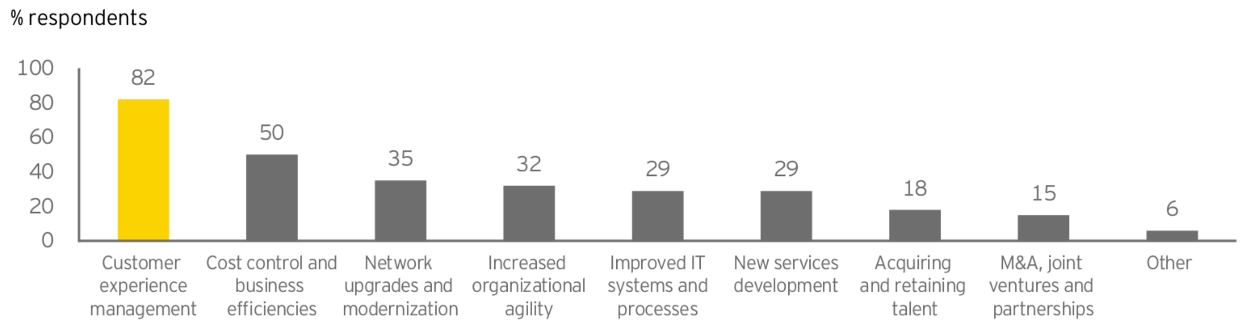

Figure 1: Leading operator initiatives to improve the customer experience [1]

Carrying higher traffic volumes, with a heavier focus on customer experience results in increased network complexity, which in turn drives higher operational costs; however this is in contradiction with the recurrent financial pressures to reduce OPEX under a stagnating ARPU and declining revenue scenario.

To solve this dichotomy, Telco Operators have typically approached the OPEX reduction problem in two different ways:

- Network Operations outsourced by contract to other industry agents through Managed Services, such as OEMs and other third parties. Despite some success stories, such outsourced activities have proven difficult to manage in the long run, both from a contractual and operational point of view, often resulting in compromised customer experience and limited cost savings.

- Network Sharing agreements with other competitors to reduce the operational cost through different schemes, sharing infrastructure or spectrum assets. Again, operational benefits have come with their own challenges such as coordination of different company cultures, priorities and strategies, as well as operational complexities derived from the reduced control of the network elements given away to the partner operator.

There is, however another option to increase operational efficiency which is becoming increasingly clear to technology leaders: the Automation of inefficient, manual processes. The development of Big Data technologies and the rise of Artificial Intelligence in an industrial scale are key enablers for this new revolution.

Automation fueled by AI

Unlike preceding technology transformations, the Automation of Network Operational processes is very subtle and does not require major investment. At the same time, it doesn't provide huge savings with single ticket items, but rather, it provides significant cumulative savings when a multitude of automated “use cases” are deployed.

Some operators will have a hard time getting started on this concept. In discussions with executives of a tier-1 operator not long ago, they mentioned that they wouldn’t consider a cost-reduction measure that didn’t provide savings over $100M. While there are some heavy processes that could deliver those savings levels, the majority of the savings will come from smaller, use cases, which many times are hidden costs to the operators. This is indeed a “revolution by incremental change”, and in TUPL we envision that Network Automation driven by AI will eventually achieve operational efficiency gains on the orders of magnitude above any other existing option, while maintaining the operations know-how and control fully within the company´s oversight. Indeed, some of our most advanced customers are targeting automation levels on the order of 80% in their network operations in the not-so-distant future.

Artificial Intelligence will be an enabler technology for this change, however it will not be the only one. The benefit of AI for Telecom is that, when properly used, it can help develop use cases very quickly, unlike traditional software development approaches. In addition to this, AI technologies can capture complex, “grey area” decisions which are often made by engineers in their day to day work, and evolve over time.

We are also conscious that the introduction of AI will come with its own challenges and that a sensible approach is required to make it a success story within Telecom Operators. In the following sections, we will describe our vision and the case for AI with particular use cases.

What is Artificial Intelligence?

Artificial Intelligence (AI) is a generic term involving a number of techniques aimed at achieving human intelligence capabilities such as learning, planning and reasoning in machines.

AI has been in the spotlight several times over the last fifty years. There have been cycles of excitement about the potential of AI upon new breakthroughs followed by the so called “AI winters” after practical limitations were hit.

During the last decade there have been a series of developments that have brought AI technologies back to the spotlight:

- Massive increase in the amount of available datasets.
- Development of Big Data technologies by IT and Internet industry leaders.
- Further breakthroughs in key areas of AI, such as Machine Learning and Deep Learning.
- Availability of low-cost, massive computational resources

Machine learning is essentially a form of applied statistics that intends to develop algorithms that are able to learn from data to perform specific tasks such as classification, recommendation, regression, transcription (i.e. speech recognition) or translation.

Most machine learning algorithms can be classified within two different categories: Supervised and unsupervised learning. Some of the most famous Machine Learning techniques are Feature Learning and more specifically Deep Learning, thanks to the remarkable breakthroughs achieved in the last few years [3]

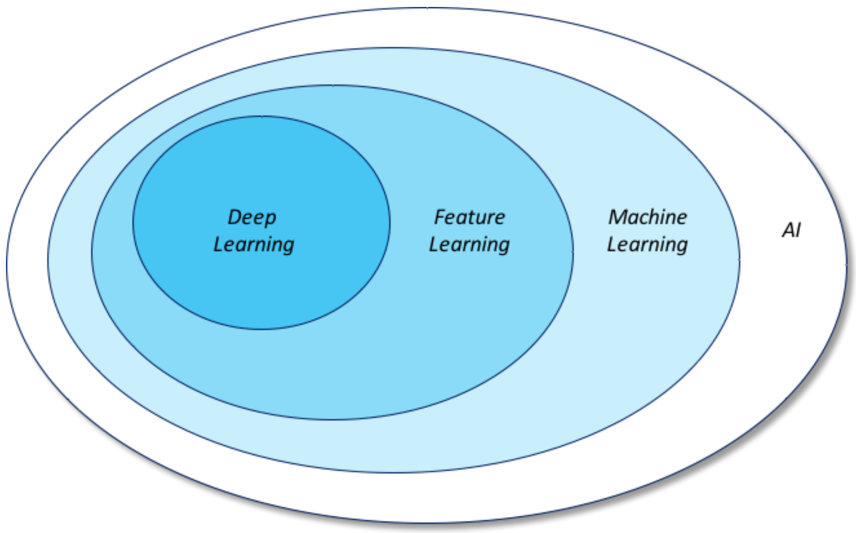

Figure 2: Venn Diagram of AI technologies

Every day, ML is used across a variety of applications, ranging from image recognition for the self-driven car or medical industry, to speech recognition, or to predicting diseases many years before symptoms develop [4] . Generally, the goal of ML is to enable a system to learn from the past or present and to use that knowledge to make predictions or decisions regarding unknown future events. Under this definition, ML can be seen as a tool that can greatly help in creating smart automated applications.

Why do operators need AI now?

Let's consider the case of cellular operators. In the past five years, network traffic has grown 18-fold [2] , and one can easily make a case that network complexity has increased by at least a factor of 5x-10x with increased cell sites, technologies and frequency bands. Managing the networks in these conditions have become a titanic task for any engineer, no matter how good he/she is.

Take for example, the case of troubleshooting customer complaints about VoLTE dropped calls. Customer care personnel must first reverse engineer the sequence of events by collecting traces and metrics on every network element/protocol plane. Then they must attempt to retrace the most likely root cause, which could be a problem with the terminal, the network element or even the traffic or coverage conditions for that user. Sometimes these investigations may take several days. Not only does this highly manual exercise require skillful engineers, but it is also a complex, costly and lengthy exercise. It is also worth noting that when solutions are found, they remain restricted to a given customer, at a given time, and a specific location: they remain local solutions that could be symptomatic of a global problem. Furthermore, the experience learned by the engineers that have conducted the troubleshooting effort doesn't get shared with the rest of the organization. In other words, the collected knowledge is often a victim of the inherent silos.

There is no other scalable way other than automation to continue to survive both the sheer operational complexity, and to serve customers better and faster. However, capturing these complex decision making processes is not an easy thing to do unless one considers the use of AI technologies, in particular Machine Learning. ML algorithms are a good fit for engineering decisions because there are many factors to take into account, and the decisions are never black or white. A good combination of unsupervised and supervised machine learning techniques can save tremendous amounts of time for engineers that are teaching the system how to behave, and then these algorithms can take over the task while the engineers concentrate on other matters.

What then is the road towards Telecom oriented AI based production systems? When it comes to applying AI or ML algorithms to operators, AI in itself is not straightforward. Deploying AI for telecoms requires a lot more components and skills than predicting shopping recommendations, or identifying whether an image is a cat.

Operational AI for telecom operators; key considerations

Continuous, Telecom grade operations

It is important to understand the difference between a research exercise and an operational system that leverages AI. Many AI companies or R&D teams perform standalone exercises that provide useful information at one point, but then the algorithms are forgotten or not properly maintained. While it is relatively simple to do "one-off" analysis based on an offline set of data, build and train a model, and extract results, doing that on an ongoing, failure-proof, telecom-grade level operation is not as straight forward.

Furthermore, achieving aggressive levels of automation will require a large number of automation utilities or use cases running together. Coordinating the execution and the exchange of information between them is an important aspect to be considered.

Data processing challenges

Everybody knows the expression "garbage in, garbage out", and in the case of AI logic this is even more relevant. Telecom data is very "noisy": there is too much data, changing over time, and often difficult to understand. Extracting value out of these feeds requires deep domain expertise, and the typical data scientist would be at a loss. Not to mention the sheer volume of some of these feeds, such as protocol traces or DPI records, which make it very hard to process in a continuous, online fashion.

In addition to this, getting access to the data itself is not easy. Data feeds are governed by different departments, and orchestrating the right permissions is typically an ordeal. In many cases, the data feeds may contain sensitive data, such as CPNI or PII, which would require the use of special security handling methods.

The tasks around data acquisition, cleanup and transformation that are required in order to feed an ML model can easily take 60-80% of the overall effort of building an AI-based solution [5] . It is important to use a set of utilities to simplify and minimize the effort in this process.

Creation of a Digital Knowledge Base

Despite all the technological advances in the Telecom sector, the fact is that the information required to develop and train AI models is not available today in digital formats. A crucial aspect to consider is how to transfer the knowledge that resides in engineers' brains to a computer network.

This Digital Knowledge base has two components: the "things to look for", or features, and the "decisions", which is a combination of models and training samples.

It is important to consider how to facilitate the generation of this Digital Knowledge base, and to do it in a way that is a cumulative and centralized effort. This will enable a continuous improvement in the company processes and will enable interesting interactions between various engineering teams. When provided with the right tools, engineers will shift their focus from day to day operations to maintaining and enhancing this knowledge base.

The creation of these mechanisms will enable operators to capitalize on their operational costs, by storing all the engineering knowledge into a digital format, which then is used to guide the various automation utilities that are used to maintain the networks.

Need for Transparent Control

Communication networks are an essential asset of today's societies. It is therefore critical to be able not only to automate the systems, but to have a way to look into the history of the network operation decisions and actions, in order to assess whether the system is making the right choices, to understand why those choices were made, and to continue training and steering the decisions. This is even more important in cases when parts of the network become unavailable, because of equipment failures, natural catastrophes, or even become the victim of malicious attacks. The need to avoid a "black box" approach puts some restrictions on some of the ML techniques that can be leveraged, at least for certain critical decisions. For example, at this point it is quite hard to understand why Deep Learning algorithms make their decisions, however there is quite a bit of active research on this topic and it is likely that this won't be a problem in the near future.

In summary, all these considerations call for a framework with the right set of AI utilities, an "AI Engine" that will facilitate the development and deployment of a multitude of use cases, the collection and correlation of the data, the creation and training of the models, and the operation at telecom-grade levels of security and availability.

Figure 3 shows that TUPL's experience in this field has helped shape a vision of an AI Engine which is not only 100% aligned with what has recently been presented by thought leaders in the AI world [6] , but one that also encompasses additional components specifically needed to meet the demanding level of quality required by telecom operations, such as security and privacy control management, continuous operations in a high availability environment, and app and user management utilities.

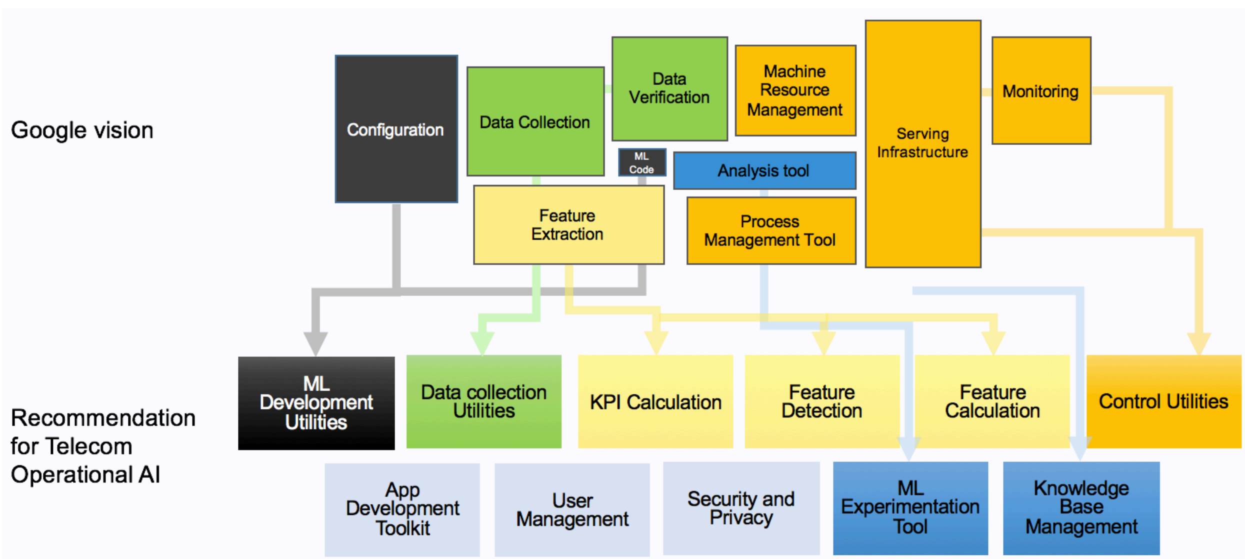

Figure 3: AI platform key components. Top: Google's vision, Bottom: TUPL's recommendation for network operations (ML: Machine Learning).

AI case study: automating technical customer care

Many companies have recognized Intelligent Process Automation (IPA) as a tool that promises to solve the complex business exercise of keeping revenue growth at a stable pace while operating their business in macroeconomic conditions that put a lot of pressure on profit margins. IPA can be defined in a variety of ways, and perhaps Mc Kinsey's recent article [7] , best summarizes how IPA should be considered. In a nutshell, IPA mimics activities carried out by humans and, over time, learns to do them better. Its application in other industries have yielded automation levels of 50-70% with 20-35% annual run rate cost efficiencies. There is little data available in the press to compare these achievements with their application to the telecom industry. This paper introduces real life results for applying IPA to telecom operations, in particular, to the process of handling technical customer care requests which, upon the best understanding of the author, is first of its kind.

For most operators today, Customer Care is a long process, where complex customer issues may involve hundreds of engineers, and several days to identify a root cause and to apply network fixes. This often leaves the customer dissatisfied as the personnel are unable to address the customer issue in a timely manner. It also generates customer frustration since in many cases the responses are unsatisfactory and they don't solve the customer problem fast enough. Further, for many operators, the resolution of these "tickets" is a hidden cost that is sometimes hard to quantify. Engineers will be spending their time doing these tasks, but no real value is extracted from it, and often the results are mediocre because there is not sufficient time to do a proper investigation.

TUPL's Automatic Customer Care Resolution (ACCR) use case was developed to target this specific problem. Built using a combination of data inputs, that range from network node metrics to customer records, it applies a combination of supervised learning models to find the most likely root cause for a given customer complaint. The utilities provided by TUPL's AI Engine facilitate the training of the data sets, as well as the continuous feedback to keep the tool operating at the proper accuracy levels.

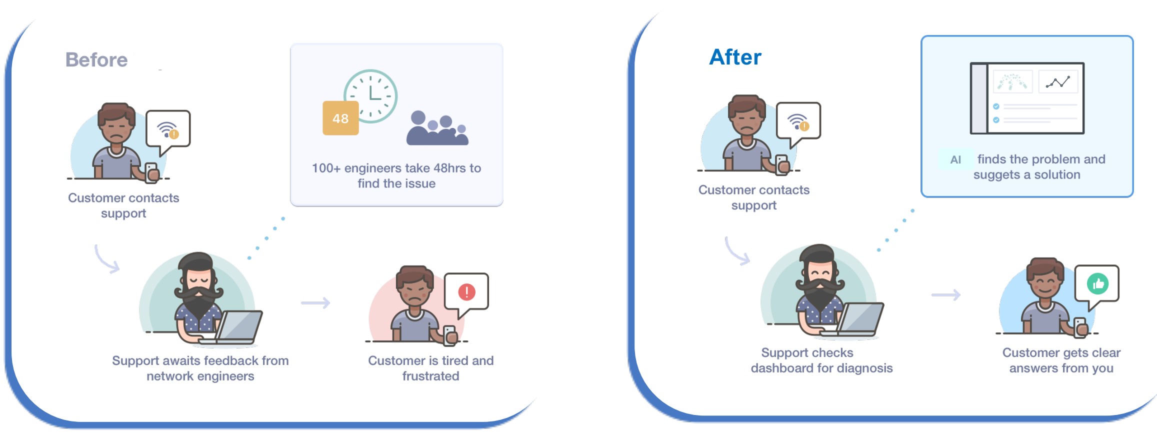

Figure 4: Transformation of customer care operations and massive improvement in user experience. Left: User experience before deployment of AI platform, Right: user experience after deployment of AI platform.

With TUPL's automatic customer care solution, Figure 4 right shows that CC personnel no longer have to consult and depend upon an army of engineers to identify the customer complaint root causes. Through a simple User Interface (UI), the tool delivers a comprehensive and simple resolution which indicates the most likely root cause. Functionality does not stop there: it also comes up with urgent escalations when needed, and suggestions for network fixes. From a business perspective, the presented automation use case has not only dramatically improved the accuracy of root cause analysis, but it has also reduced the number of engineers required to address a customer issue, resulting in improved value while lowering operating cost.

We reproduce Figure 5 with the initial results of the automation use case for technical customer care.

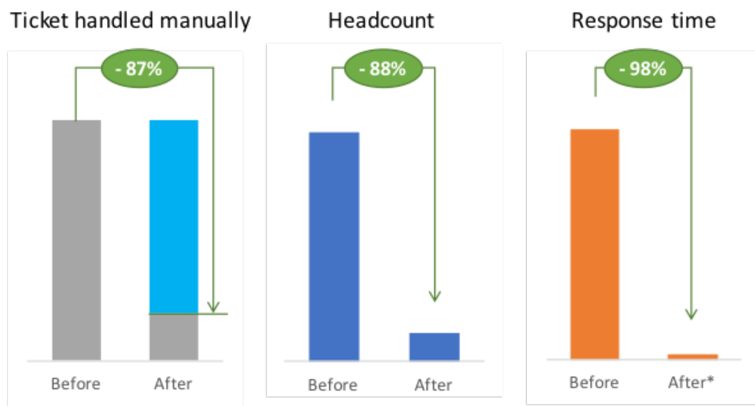

Figure 5: Efficiency of TUPL's automatic customer complaint resolution AI platform at a Tier 1 mobile-operator. "After" refers to the status after deployment of AI platform to handle customer tickets resolution.

Data collected from this use case show that TUPL's solution produced the following results:

- Near 90% of customer care tasks can be automated (Figure 5-left)
- A factor of 4 reduction in the "undetermined" root causes
- Nearly 100 times faster response to customer issues has been reported

This featured solution has not only completely changed the way Customer Care (CC) is run by conventional teams, it has also raised the user experience to a level that could not easily be achieved in traditional organizations.

One must remember that the case reported here is only one example (albeit a powerful one) for operations automation by AI. Once the first solution is deployed, it will be even easier for the subsequent ones as the correlating and learning engine (i.e. AI Engine) is already in place. Furthermore, the digital knowledge base gathered during this phase facilitates the development of the next level solution, Proactive Care, which is focused on identifying and resolving customer issues before the complaint occurs.

Reference documents:

[1] Ernst & Young, Global telecommunications study: navigating the road to 2020 https://webforms.ey.com/Publication/vwLUAssets/ey-global-telecommunications-study-navigating-the-road-to-2020/$FILE/ey-global-telecommunications-study-navigating-the-road-to-2020.pdf

[2] Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2016–2021 White Paper https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/mobile-white-paper-c11-520862.html

[3] The Guardian, "World's best Go player flummoxed by Google's 'godlike' AlphaGo AI", May 23rd 2017, https://www.theguardian.com/technology/2017/may/23/alphago-google-ai-beats-ke-jie-china-go

[4] The Douglas Mental Health University Institute's Translational Neuroimaging Laboratory at McGill Montreal, Canada, "Artificial intelligence predicts dementia before onset of symptoms"August 22nd 2017.

[5] Kai Wähner, Data Preprocessing vs. Data Wrangling in Machine Learning Projects, March 2017

[6] D. Sculley et. al. (Google inc.), "Hidden Technical Debt in Machine Learning Systems", NIPS'15 Proceedings of the 28th International Conference on Neural Information Processing Systems, 2015, p2494—2502

[7] McKinsey & Company, "Intelligent process automation: The engine at the core of the next-generation operating model", March 2017, **http://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/intelligent-process-automation-the-engine-at-the-core-of-the-next-generation-operating-model**