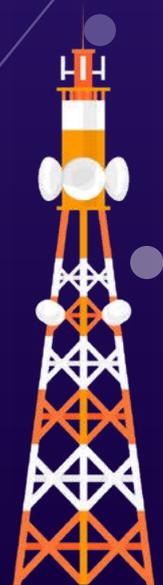


AI with
reinforcement
learning:
**the secret sauce
for 5G success**



The demands of a 5G world

There is no doubt, the adoption of 5G is going to leave all other network technology upgrades in the proverbial dust. With the promise of 5G for 10 to 100 times faster data rates, it is expected that existing services will receive an unprecedented upgrade in performance.

The more dramatic tectonic shift of 5G, however, is in how it is poised to power life-changing new services, including for telehealth, autonomous transportation, and many more.

But operators are faced with many challenges in delivering these exciting new services, with the 5G network being **extremely complex** to plan design, launch, and operate.

Moreover, there is a need for **real-time decisioning and response** in order to effectively allot the network resources required for delivering innovative services.

Real time is so important because the 5G network is in **constant flux**, with **consumption varying** by the use case and the users, requiring the RF environment to be continually adjusted.

To illustrate the complexity and how it impacts the ability (or inability) to optimize the performance of a complex 5G network in real time, let's take a look at the Massive MIMO antenna.

And this is where we can see the introduction of complexity. Namely, passive antennas have only one beam and require only one parameter to be configured, i.e., antenna tilt. But with active (Massive MIMO) antennas **many parameters** must be configured, including beam number/forming, tilt, azimuth, and vertical and horizontal aperture. Moreover, the **number of beam patterns** can reach 10,000 and even more.

The Massive MIMO antenna also introduces greater complexity over its legacy predecessor with regards to the following:



Each 5G cell has its own **unique physical attributes** and coverage objectives that need to be planned for.



The relevant **geography**, including buildings, stadiums, theaters, roads, and others, is constantly changing and needs to be addressed on an ongoing basis.



User and **traffic distribution** is also continually shifting.

The active (Massive MIMO) antenna challenge

The passive antenna, which is commonly used in, UMTS, and in LTE networks, has been replaced in 5G networks by the active antennas (e.g. Massive MIMO), with its large number of radiating dipoles.

This change has been enacted since Massive MIMO antennas enable operators to drive more traffic to bandwidth thirsty 5G applications, as well as further to their ability to boost channel capacity and to act as the infrastructure for beamforming transmission.

Moreover, active (Massive MIMO) antennas in a 5G network offer increased flexibility in controlling the way RF signals are propagated, with their ability to have the number of beams, beam boresights, vertical beam widths, and horizontal beam changed as needed.

This means that the profoundly **complex configuration** of the active (Massive MIMO) antenna is not only challenging **at the outset**, but also during **ongoing operations**. And all this needs to be executed continually in **real time** to accommodate an ever-changing environment.

As such, the Massive MIMO antenna mandates configuration and ongoing operation that is **automated, dynamic, intelligent, and real-time**.

This is where AI comes in.

MIMO Beamforming



How artificial intelligence can help

“Salt and pepper. Day and night. Fred and Ginger. Some pairings create an exquisite experience that’s simply not otherwise imaginable. The same is true of artificial intelligence (AI) and 5G networks.”

CIO.com

AI is not new to operators. Projects abound worldwide, whose aim is to leverage the technology for extracting strategic insights from data to improve the customer experience, enhance operational efficiencies, reduce capital expenditure, optimize network performance, enable hyper-automation, and to identify new revenue generating opportunities.

And when it comes to the most commonly used AI learning techniques, it’s all about supervised and unsupervised learning.

Supervised learning

Supervised learning involves the use of labeled datasets to train algorithms to classify data and predict outcomes accurately. By learning from pre-labeled training data, the algorithm can make predictions for unforeseen data. It is often likened to learning in the presence of a teacher or supervisor (hence, ‘supervised’ learning), and its intent is to identify patterns and correlations that will be applied to analytics processes.

Unsupervised learning

Unsupervised learning, also known as cluster analysis, doesn’t rely on supervising the model. Instead, the model operates independently to discover information from massive amounts of unlabeled data.

Understanding the meaning behind unlabeled data requires algorithms that perform classifications and clustering as based on the patterns they find, through an iterative process that requires no human intervention.

Advantages and disadvantages

AI solutions for operators that are driven by supervised and unsupervised learning do have their advantages. For example, supervised learning is very helpful in classification problems and can help to predict a target numeric value once given a set of features called predictors.

Unsupervised learning is helpful in finding patterns in data that are not discernable through other methods.

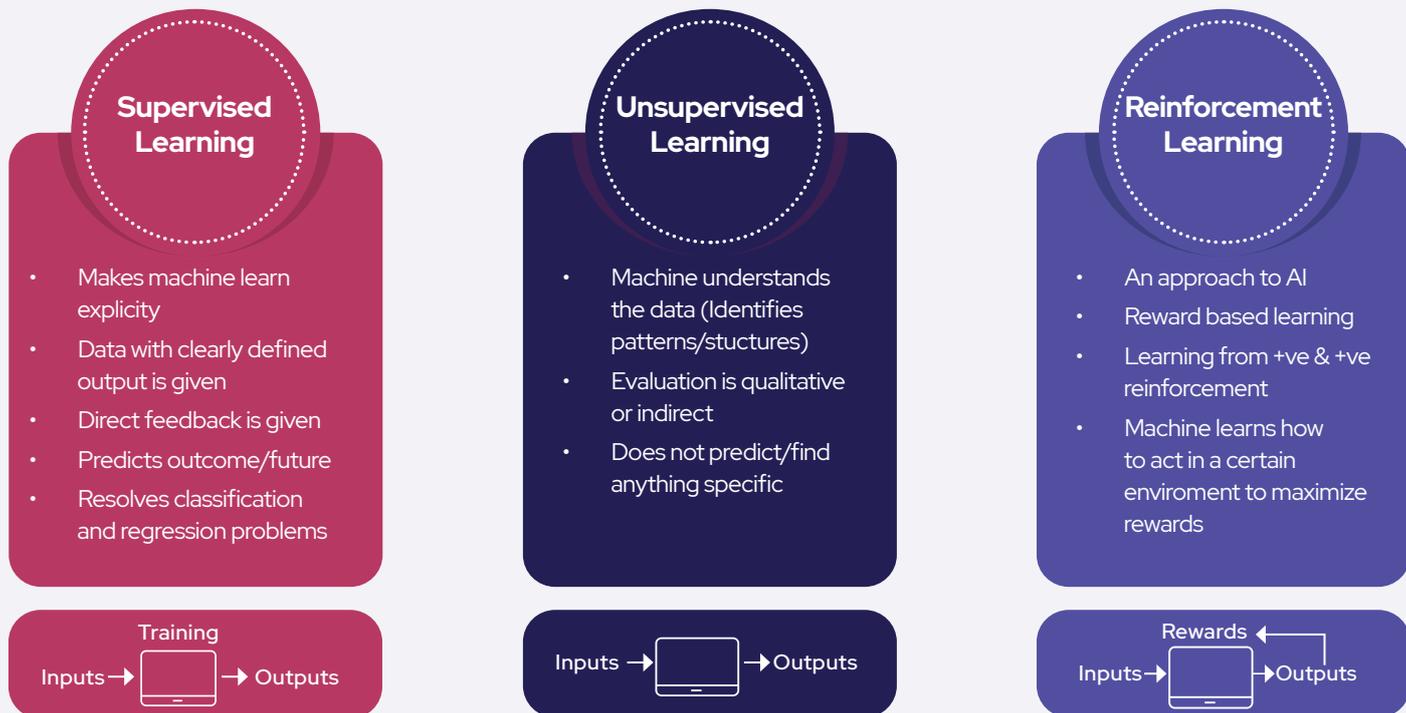
However, the great disadvantage of both is that **neither is designed for delivering real-time insights**. While supervised learning delivers accurate results, these cannot be applied for real time insights.

And while unsupervised learning can deliver real-time insights, there is a high risk that the results will not be sufficiently accurate.

So, we can see that neither is applicable for meeting the needs of real-time decisioning and resource allocation in a 5G network.

In comes reinforcement learning

The good news is that there is another kind of learning that can handle the real-time needs of 5G network performance optimization. This is reinforcement learning (RL).



Reinforcement Learning is different, because in this case the machine, just like humans, learns in real time and does so through trial and error.

The learning system also known as the agent, independently learns what is the best strategy, known as a policy, for attaining the most positive reward (outcome) over time.

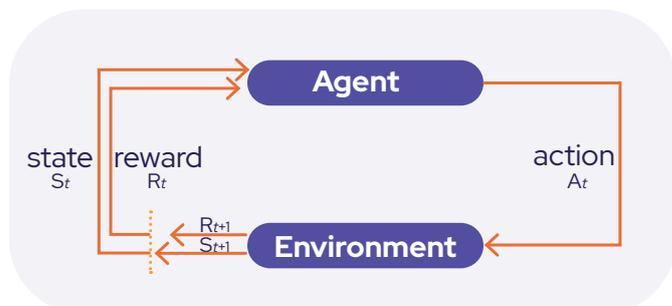
This process typically entails six steps:

1. Observation
2. Selecting an action using the best strategy
3. Taking action
4. Attaining positive or negative rewards
5. Analyzing the rewards and updating the strategy
6. Repeating the process until the optimal policy is obtained

Ultimately, a sequence of successful decisions, which garner the most positive signal rewards, will result in a process being reinforced, hence the name for this type of learning.

Reinforcement learning as the scheduling agent

In the 5G environment, reinforcement learning can serve a scheduling agent that continuously samples the network in real time. And through an ongoing cycle of trial and error, i.e., observation–decisioning–analyzing outcomes, it enables continual and real-time network performance optimization.



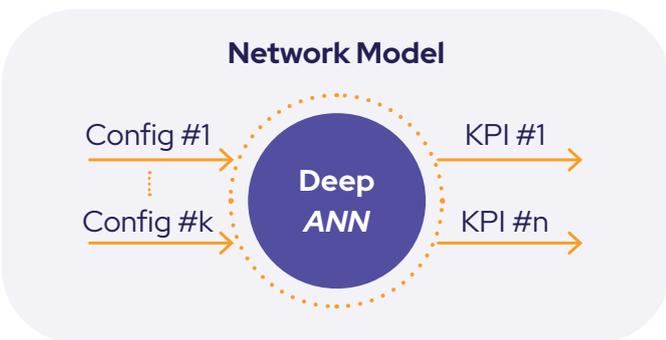
The importance of gradual implementation

Though, with all of its powerful capabilities, if reinforcement learning does all of its trial-and-error learning in a production environment, it will impact quality of experience (QoE). For, during the learning process, RL can impact the network's configurations, and this may result in a significant degradation of network performance and deterioration of the user experience.

The key to overcoming the challenge is **gradual implementation**, where initial learning takes place **offline in a lab environment**, and only thereafter continuous learning is moved to production.

To see how this plays out, let's take a look back at the Massive MIMO antenna.

Let's say we want to predict the impact of re-configuration on the network. To do this we would need to develop an auto-configuration AI model that is driven by big data and deep artificial neural network (ANN) learning (Deep ANN is an ANN with multiple layers between the input and output layers).



In **phase one** of learning, the training should be performed offline in a lab environment that has been developed for implementing deep ANN with supervised learning.

Doing so enables us to learn about how changing the Massive MIMO configuration for coverage, traffic, interference, throughput, load, spectral efficiency, and more would impact cost and what is the optimal configuration at the cell level.

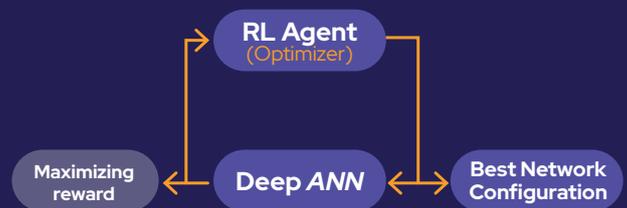
It will also enable us to predict how any change in the configuration of one cell will impact the relevant neighboring cells, not just the cell itself.

Once we know all this, we can now go into **phase two** and leverage reinforcement learning for leveraging the results of the deep ANN with supervised learning offline as well as for executing required optimizations on an ongoing basis.

The RL agent senses the changes in the radio environment as well as in the various data services, both offline and in live production. And, according to the numerical reward function that it determined (per the process outlined above), it "decides" what the optimal Massive MIMO configuration would be per change.

In short, we can see the best way to assure no impact of Reinforcement Learning on QoE is to:

- Have initial learning performed by deep ANN with supervised learning.
- The results are fed to RL still offline, to observe permutations and analyze the impact of configurations (produced by deep ANN + supervised learning) in the lab environment.
- RL can then refine predictions for the optimal configuration and implement in the live environment.
- RL executes ongoing learning in the live environment, making optimizations when necessary.



This RL-powered model provides MNOs many benefits, including:



How Cellwize can help

Cellwize CHIME applies reinforcement learning to automated orchestration and optimization of the 5G RAN, helping operators extract the most value out of AI for optimizing 5G network performance.

Among the reinforcement-learning powered use cases Cellwize supports are the following, with new ones continually being introduced;

Use Case	Description	The Cellwize approach
Massive MIMO configuration optimization	Continuous learning of all the best configurations for every parameter on every antenna, and optimization of the configuration in real time to enable a high performing 5G network	A combination of deep learning and reinforcement learning to produce a Massive MIMO configuration optimizer. The Cellwize Optimizer solution can be configured to optimize a wide variety of pre-specified cost metrics, such as those that involve capacity/coverage tradeoffs.
Auto-tuning cluster parameters	Leveraging reinforcement learning for automated policy enforcement	Leveraging auto-classification for labelling and clustering cells, and then applying reinforcement learning to auto-tune the parameters of the cells in the clusters.
Recommendations and federated learning	Leveraging reinforcement learning to drive recommendations and federated learning for model training across multiple decentralized edge devices or servers, without the need for exchanging device data. Furthermore, federated learning can also be used to reduce the need for transporting data from radio nodes to other locations, which helps to address data privacy, access, and security issues.	Leveraging deep reinforcement learning and combining artificial neural networks with a reinforcement learning architecture that enables software-defined agents to learn the best actions. The deep RL agent is trained and re-trained in distributed machines. Then, the trained model interactively updates a centralized model. The idea is to avoid the sharing of user data and update the machine learning model in a distributed manner.
Antenna tilting	Leveraging reinforcement learning with an agent running in the core that is trained to dynamically control the electrical tilt of multiple base stations. This process is performed jointly across multiple BTSs to improve a cell's signal quality and to reduce the interference on neighboring cells.	Leveraging RL for antenna tilt angle optimization for macro-cells in a heterogeneous cellular network (HetNet). The RL algorithm is applied in two phases. During the first phase the agent learns the behavior of the neighboring cells for each state-action in an offline mode. During the second phase, a deep neural network learns the user equipment's locations. This way the 'reward' can be estimated without taking any action.

With CHIME, operators get all this and more. And, ultimately they get a smarter network so they can accelerate the deployment of their 5G network and optimize its performance easily, safely, and without risk. **Only rewards.**





About Cellwize

Cellwize is all about enabling the networks of the future today. With CHIME, our cloudified and AI-driven RAN automation and orchestration platform, we enable mobile network operators (MNOs) to accelerate 5G network deployment and go-to-market, as well as the ROI on their network investments. Even in the most complex and dynamic of network environments, CHIME enables operators to connect to any application and any vendor, as well as co-create on top of the platform, delivering unprecedented ease, speed, and agility. With the future of 5G already here, Chime is helping leading MNOs all over the world to launch and leverage their next generation networks and face the future with confidence.

To learn how reinforcement learning can help you expedite the journey to 5G, we invite you to reach out to us [here](#)

