AI opportunities in 6G Layer 2

White paper

With advances in end-user devices and other technologies, the demand for immersive experiences that seamlessly blend the digital and physical worlds will be significant by the end of the decade. To support these, 6G networks will need to deliver unprecedented capacity, low latency, energy efficiency, and cognitive capabilities to manage vast radio resources. The potential of artificial intelligence (AI) to enhance the physical layer and reach these goals has been demonstrated in previous works. In this white paper, we explore the frontiers of machine learning (ML) to improve features beyond the physical layer (PHY) and further enhance the native AI air interface (AI-AI) envisioned for 6G. We survey recent research on AI-driven functions such as resource allocation, random access, adaptive modulation and coding (AMC), power control, protocol learning, channel state information (CSI) reporting, hybrid automatic repeat request (HARQ), and multi-RAT spectrum sharing (MRSS). We contend that while the fundamental duties of the 6G wireless medium access control (MAC) will remain consistent with prior generations, the integration of ML methodologies will instigate transformative changes across multiple MAC domains. This infusion of ML-driven strategies will open new avenues for development and maintenance and undergird future redesign efforts.

Contents

Introduction	3
Resource allocation	5
Deep scheduling of data channels	5
Random access	6
Link adaptation	7
ML-based adaptive modulation and coding	7
ML-based power control	7
Signaling overhead reduction	9
Protocol learning	9
Reporting of channel state information (CSI)	9
Multi-RAT spectrum sharing (MRSS)	11
Conclusion	11
Abbreviations	12
References	14

Introduction

The upcoming generation of mobile networks, slated for deployment at the end of this decade, will revolutionize human communications by unifying the digital and physical worlds into a scalable metaverse. New technologies such as extreme massive multiple input multiple output (MIMO), joint communications and sensing (JCAS), native AI-AI [1], and cognitive networks are being researched for 6G to fulfill these requirements (see Figure 1). At the standardization level, the journey towards 6G is also starting in 3GPP Rel-19 [2]. AI techniques hold promise for optimizing the physical layer of the air interface, simplifying complex optimization problems, and enhancing performance of known procedures such as joint channel estimation, detection, nonlinear distortion removal, and even neural receivers. Similarly, AI will play a crucial role in other layers of the protocol stack, addressing the increased heterogeneity of 6G use-cases. Indeed, Nokia has recently proposed [12] a novel protocol stack design, where the stack is split into a slow and a fast stack, to avoid low-speed optimizations damaging the high-speed user-plane. This flexible and scalable architecture would offer a high degree of customization that can be leveraged with AI methods to train protocol stacks (or parts thereof) for custom needs.

Figure 1. Key technologies in 6G

In modern mobile communication systems, the MAC has evolved beyond its basic role of wireless channel access control. Jointly with the upper PHY and other lower layers (see Figure 2), they now encompass functions such as:

- Scheduling and logical channel multiplexing, including dynamic spectrum sharing
- Link adaptation (coding rate, modulation, transmit power, timing advance)
- HARQ management
- Link and device status reporting (channel state feedback, buffer status reporting (BSR), power headroom reporting (PHR), link and beam failure)
- Energy efficient features (discontinuous transmission and reception).

Figure 2. Main functionalities of the 6G layer two (L2)

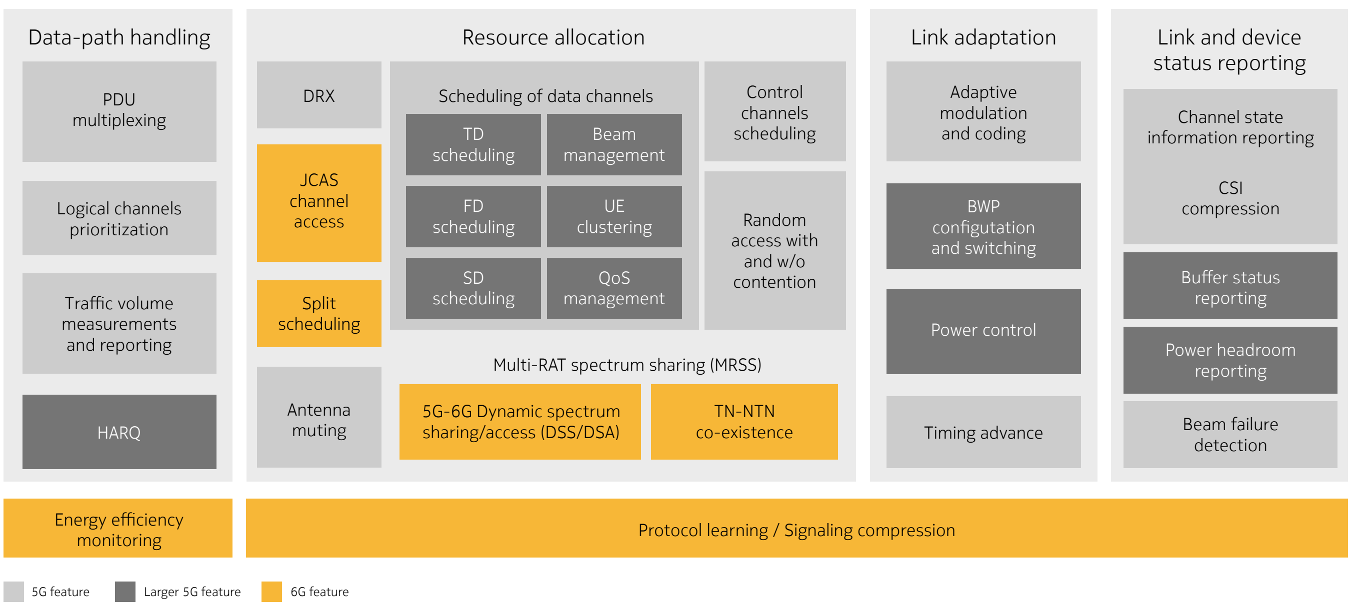

As these functions grow increasingly elaborate, ML techniques hold the potential to comprehend and optimize them effectively. To exemplify the escalating complexity, let's consider the packet scheduler — a quintessential MAC function. In massive MIMO macro systems, the scheduling function involves MIMO pairing to co-schedule users across multiple spatial layers, utilizing the same time/frequency resources. 5G already supports a limited version of multi-user MIMO (MU-MIMO). The major challenge, however, remains to scale the number of co-scheduled users even faster due to smaller time slots.

In 6G, the scheduler will face additional challenges, including managing network slicing with increased carrier bandwidth (100–400 MHz), accommodating diverse QoS requirements, and co-scheduling 5G and 6G devices on shared carriers. The intricate dependencies between these responsibilities will challenge AI techniques to unveil novel packet scheduling methods. Rel-18 3GPP studies already explore AI applications such as beam management and channel state feedback compression for 5G-Advanced, and some, such as beam management, will be standardized in Rel-19 [2]. In addition, radio MLOps, as recently proposed by Nokia [3], will certainly be needed to enable all these functionalities in a scalable manner.

Resource allocation

Deep scheduling of data channels

Traditional packet schedulers are structured in sub-modules for time, frequency and spatial-domain scheduling. In the first sub-module, a time domain (TD) scheduler identifies high-priority user equipment (UE) for scheduling on each transmission time interval (TTI). Subsequently, a frequency domain (FD) scheduler maps frequency resources to the selected UEs. This may involve beam selection and MU-MIMO pairing.

Figure 3: Performance gains (left) achieved by a DDQN-based scheduler over a PF baseline under various traffic. The learning curve on the right illustrates learning progress during training.

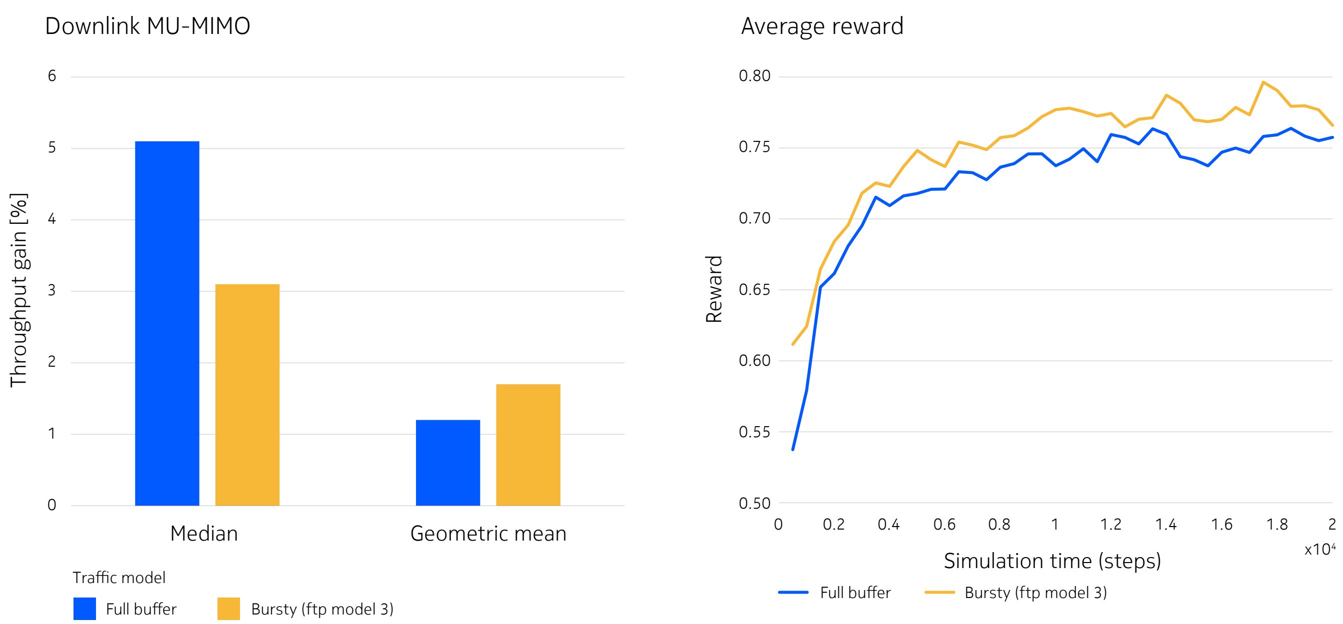

To illustrate deep RL-based scheduling, we trained a downlink MU-MIMO scheduler using a double-deep Q-network (DDQN). Imitation learning techniques were used to avoid performance-damaging random exploration, where the DDQN agent was rewarded when its decisions matched those of an expert Proportional Fair (PF) scheduler. The resulting deep scheduler learned to leverage frequency selectivity and dynamic frequency-hopping patterns (see Figure 3). Since training can be done offline in simulators, the agent can be exposed to a vast amount of radio challenges and learn from them. Furthermore, ML models also offer the advantage of fast real-time scheduling decisions with a constant number of operations during the forward pass. Nokia is currently conducting additional research into RL techniques for radio resource scheduling, such as soft actor critic (SAC) and proximal policy optimization (PPO), including realistic assessments of each method's training to performance tradeoff.

Random access

Random access is the first uplink transmit procedure executed by UEs to achieve connectivity. The number of random-access channel (RACH) resources is limited, however, and the lack of initial coordination by the network makes random access a contention-based process by necessity.

Conventional threshold-based techniques for the detection of random access (RA) preambles leverage the correlation properties of Zadoff-Chu sequences, which are used as orthogonal preambles. This approach, however, leads to ‘preamble collisions’ whenever two or more nodes choose the same preamble. In massive machine-type communications (mMTC) networks, the chances of several UEs choosing the same preamble are just too high. Another problem of this approach is false peaks (i.e., preambles) detection, which worsens under interference-limited conditions.

In 5G networks, low-latency connectivity was introduced and will be further enhanced in 6G. With a larger number of devices using low-latency services, situations with low coverage may lead to increased contention for limited RACH resources. Rapid resolution of these frequent contentions will be crucial to ensure the expected low latencies in 6G services.

ML is well-suited for classification tasks like preamble detection in RA. In this context, solutions based on decision-tree classification (DTC), naive Bayes, K-nearest neighbor, and Bagged decision tree ensembles for the purpose of RA preamble detection have been proposed. Experimental results show that ensemble methods achieve similar false peak detection rates as the Zadoff-Chu sequence baseline, with faster inference speed but increased training complexity, particularly under low signal-to-interference-and-noise ratio (SINR) conditions.

Overall, ML preamble and timing advance (TA) value detection is promising to leverage non-orthogonal preambles and reduce RA latency in future massive networks. Incidentally, networks with large numbers of devices are also the ones that can produce sufficiently large training datasets to train such models, thus facilitating their development.

Link adaptation

ML-based adaptive modulation and coding

Adaptive modulation and coding (AMC) is an essential function. It selects the best modulation and coding scheme (MCS) considering channel estimates, data size, and system constraints. Challenges include balancing aggressive transmissions for spectral efficiency and conservative transmissions for latency and block error rate (BLER). For classic eMBB services and their first transmission target (a BLER of 10–20%), existing heuristics achieve good performance in matching BLER targets and maximizing the spectral efficiency. These currently used techniques, however, have a considerably higher impact on spectral efficiency when the BLER target is in the order of 0.001%, which is the case for ultra-reliable low latency communications (URLLC).

Recent efforts targeting URLLC applications have, therefore, introduced new ML-based AMC architectures that balance performance gains with computational and implementation complexity. Some approaches employ models, while others use function approximators. Multi-armed bandits have also been explored as an MCS selection algorithm, as well as Thompson sampling for enhancing MCS selection by tracking a low-dimensional representation of the SINR. Additionally, [5] focused on interference prediction using kernel density estimation to adjust MCS selection based on the interference's probability density function.

Function approximators-based AMC approaches primarily rely on deep neural networks. However, the lack of ground-truth MCS decisions as labels has limited the use of supervised learning. To overcome this, deep reinforcement learning (RL) has been explored for MCS selection policies in real-time scenarios [6, 7], where lightweight and efficient solutions need to make fast decisions for multiple UEs. Small-sized actor-critic and proximal policy optimization (PPO) models have been trained, achieving notable performance. Other solutions developed by Nokia build on top of the most used existing heuristic, outer loop link adaptation (OLLA), to reduce the implementation efforts. One such approach is presented in [8], where the authors introduce a differentiable computation graph to train OLLA hyper-parameters dynamically. As an alternative, [9] proposes to substitute OLLA's analytical formula with the actions of an RL agent to offset the SINR estimate before AMC. All the proposed solutions trade off computational complexity for reliability, which can boost URLLC performance in 6G systems.

ML-based power control

In 4G and 5G, uplink transmit power management includes open-loop and closed-loop components, which ensure spectral flatness, SINR maximization, and minimal UE battery drain. Open loop power control (OLPC) adjusts transmit power via two cell-specific parameters that are typically set via trial-and-error by the mobile network operator (MNO). While historically reliable, this performance is often suboptimal.

Exploring new OLPC parameter values carries connectivity risks, and the lack of closed-form expressions for network performance discourages MNOs from deviating from established values. To address these challenges, Nokia recently proposed leveraging black-box optimization techniques like Bayesian optimization with Gaussian processes (BOGP) [10], which yielded large gains in an operator-controlled trial in a large 5G network. Compared to conventional ML techniques, BOGP converges in fewer samples, reducing outage risks and ensuring safer optimization.

Specific 6G verticals may require more frequent adaptation of OLPC parameter values. Cell-specific OLPC settings could also be desirable. To address these needs, [4] presented another concept with multi-agent reinforcement learning (MARL) and centralized training with decentralized execution (CTDE) techniques to avoid multi-agent optimization issues. Neighboring base transceiver station (BTS) interference poses a challenge when using cell-specific power control. Independent ML solutions, where each BTS learns optimal parameters, face convergence issues due to non-stationarity. In naive independent learning, performance collapses as ML agents compete, hindering stable learning in neighboring nodes. Centralized or cooperative learning with neighboring base stations is required to overcome this. One demonstrated solution [4] is to let neighboring BTSs share information for joint learning and fair action selection.

In 6G, with massive antenna arrays, MU-MIMO is the default mode, and transmissions overlap in time/frequency but are separated spatially. This leads to the challenge of minimizing intra-cell interference. Increasing transmit power for one user raises uplink interference for others, reducing the effectiveness of conventional UE-specific closed loop power control (CLPC). At Nokia Bell Labs, we are exploring a tighter integration of OLPC and CLPC functions with scheduler decisions in MU-MIMO settings. The sheer complexity of this joint optimization problem demanded an ML-based approach based on convolutional neural networks that outperforms conventional solvers.

Signaling overhead reduction

Protocol learning

In 5G, MAC protocol data units (PDUs) carry a 6-bit logical channel ID (LCID), used to identify L2 control elements (CEs) like BSR or PHR. Specialized deployments, however, may not require these messages, whose processing imposes a heavy computational toll on the receiving devices (processing packet headers and other non-payload fields is a large computational burden in radio protocols). For example, in a low-traffic wireless network such as a sensor network with infrequent events, BSR exchanges between sensor UEs and the radio access network (RAN) may be unnecessary. Removing BSR would allow reduction in the size of MAC uplink PDUs. At the same time, this approach may not be suitable for high-capacity networks that are reliant on BSR signaling.

For instance, the laser cutting machine modeled in [11] produces, on average, data packets with a median size on the order of 100 B every 10 ms, as well as some larger packets, less frequently. Since this data volume is very small, L2 protocols in this scenario would incur a high header-to-payload ratio. This differs from scenarios like enhanced mobile broadband (eMBB) and small office home office (SOHO), where large data packets carrying video frames are common. Additionally, business-critical indoor factory (InF) traffic requires reliable channel access and strict QoS, while contention-based access may suffice for SOHO scenarios. Even within the same scenario, different applications have varying requirements. Current network resource management satisfies application QoS needs, but the MAC PDU structure remains largely unchanged. A more customized control plane is clearly desired, which will require a new MAC architecture, such as the one described in [12], where most of the control-plane signaling is shifted to an anchor protocol stack (APS).

A split stack would facilitate the customization of MAC protocol headers, which can be done manually or emerge via AI with protocol learning techniques. Early work in protocol learning [13, 14] employed a holistic approach to train an ML model capable of replacing a complete MAC layer. With these methods, AI agents at the BTS and UEs jointly learned a new L2 signaling and access policy. However, the large MAC signaling space rendered most of these methods impractical and hard to scale. Instead, recent research trends focus on reducing the problem size through techniques like state abstraction [15] with multi-agent proximal policy optimization (MAPPO), semantic communications [16], or goal-oriented communications, currently being studied in the EU-funded project CENTRIC [17].

Reporting of channel state information (CSI)

In 5G, the quality of radio channels is measured across multiple dimensions (time-frequency resources, spatial layers, beams, etc.). This collection of UE-side measurements, to be returned to the BTS, is generally named CSI, and it includes metrics such as channel matrices, rank indicators, and others. For instance, UEs regularly send the channel quality indicator (CQI) to the BTS for link adaptation purposes. However, a single number like the CQI cannot capture the multidimensional complexity of MIMO channel measurements. Large channel matrices are needed for this, but they need large bandwidths to be properly fed back to the BTS. The larger number of antenna elements and the larger 6G bandwidths will increase the size of these matrices even more, which is why efficient encoding mechanisms are needed.

Transformer-based autoencoders seem to outperform convolutional ones for CSI feedback compression. This is because they leverage frequency band proximity and space correlations across antenna elements in MIMO panels with significantly fewer parameters. Model size reduction is indeed essential for deploying ML-based CSI encoders on real-time hardware. To illustrate this, Figure 5 shows mean bitrate gains of up to 13% for a transformer-based CSI feedback architecture over 3GPP Rel-16 eType II codebooks. In this design, the CSI encoder sits at the UE side, while the decoder is at the BTS side. In general, any CSI encoding mechanism must be matched against current codebook-based feedback methods and yield clear gains in terms of channel matrix reconstruction error and bitrate performance. Only if unanimity is reached on the advantages of these methods will the relevant industrial bodies (e.g., 3GPP) move to standardize the necessary support for separate training of encoder and decoder, life-cycle management of ML functionality, and data collection.

Figure 4. Rank 2 MU-MIMO performance of a transformer CSI compression chain with 1.5M trainable parameters [18]. Uma SLS at fc =4 GHz, BW = 20 MHz, 21 cells and 80% indoor and 20% outdoor UE distribution. L denotes the number of combined beams reported by the UE.

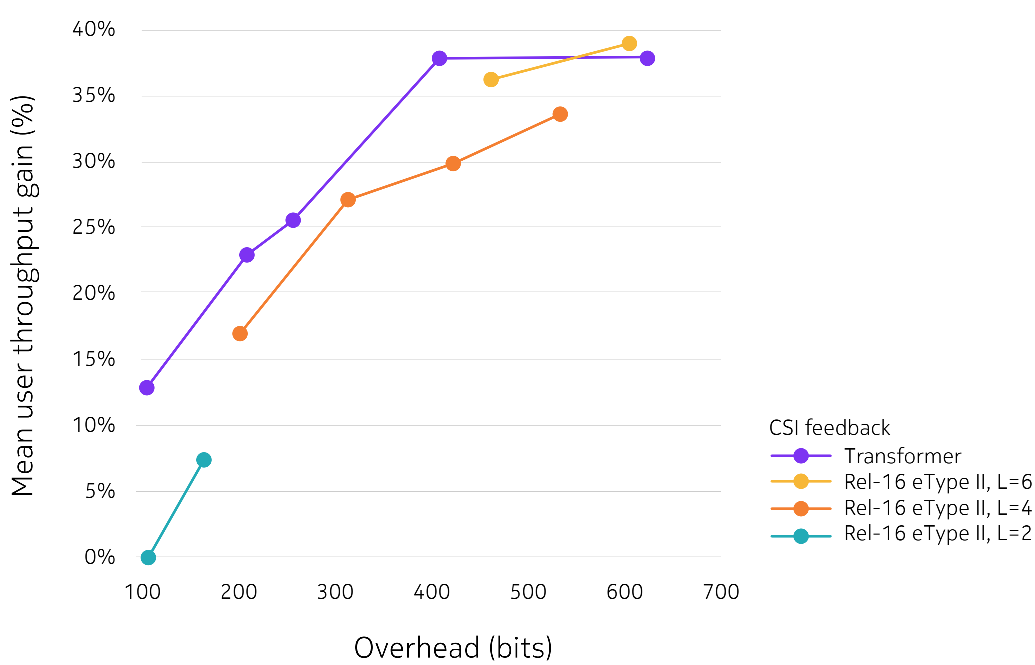

Multi-RAT spectrum sharing (MRSS)

WRC-2023 identified bands 7.125-8.4 GHz and 14.8-15.35 GHz for potential 6G usage, as they will support the large bandwidths expected of this technology (up to 400 MHz). Nonetheless, the 410–7125 MHz band (aka 3GPP FR1) will still remain as an important option for coverage and capacity. In the early 6G deployment phase, however, these bands will be crowded with legacy 5G networks. As 6G devices slowly replace 5G ones, MRSS will provide a dynamic co-existence solution to ease the transition. A key design decision in MRSS deployments will be whether to coordinate 5G and 6G spectra separately via dynamic resource splitting, or through joint 5G-6G scheduling.

The 5G-6G co-existence scenario's complexity may surpass that of 4G-5G, yet it presents increased flexibility with the beam-centric nature of both generations. Enhanced coordination between 5G and 6G schedulers will be needed to leverage these capabilities. Furthermore, 6G deployments will be more heterogeneous, involving private 6G networks sharing spectrum with public and/or private 5G networks. A single 6G design may not be suitable for all spectrum-sharing scenarios, but ML models can be trained for specific transient needs. For instance, a small private 6G network near a public 5G network could benefit from learning tailored orthogonal frequency-division multiple access (OFDMA) pilot patterns for the public network configuration. Retraining the private network as the public network evolves lets it adapt to shifting coexistence conditions.

Distributed deep MARL techniques have also been proposed for spectrum access coordination [20], which is preferable for private 6G networks lacking interfaces to nearby 5G networks. Deep classifiers have also shown [20] reliable performance in spectrum sensing tasks, accurately predicting spectrum usage. In general, MRSS is a migration feature and will not persist during the lifetime of 6G deployments. It may, nonetheless, be essential to 6G's success by facilitating the highly dynamic transition period from 5G to 6G. Further details about MRSS can be found in [19].

Conclusion

The lower layers of the radio stack, and the MAC in particular, play a crucial role in both PHY operations and L2 control-plane semantics. For this reason, the future of ML-based MAC entails using smaller, separate models that focus on specific MAC duties. Considering real-time constraints, it may not always be feasible to run all models simultaneously, thus ad-hoc protocol stack architectures, such as the recently proposed Fast Protocol Stack (FPS) [12], are needed.

But AI also opens the door to paradigm shifts in various MAC domains. Deep schedulers will revolutionize wireless product development and testing, enabling the efficient management of advanced radio systems while maintaining bounds on the real-time computational load for inference. Their application to 5G-6G spectrum sharing also seems promising. ML-based detectors for non-orthogonal preambles are also improving rapidly, offering tremendous latency reductions for random access channels. Deep policy-based RL seems the tool of choice for learning tailored AMC solutions, and Bayesian methods will optimize large network deployments with only a handful of data samples. Customized 6G protocols with low signaling overhead will emerge automatically via semantic communications. Attention-based mechanisms will enable the effective encoding of large channel matrices for CSI feedback and reduce HARQ overhead with novel transformer-based autoencoders. Although the 6G wireless MAC will retain its core responsibilities from previous generations, ML-driven approaches will introduce novel ways of development and maintenance, facilitating future redesigns. This propels the MAC layer into a realm of unprecedented possibilities and sets the stage for groundbreaking advancements in wireless communications.

Abbreviations

3GPP	3rd Generation Partnership Project
AI	Artificial intelligence
AI-AI	AI air interface
AMC	Adaptive modulation and coding
APS	Anchor protocol stack
BLER	Block error rate
BOGP	Bayesian optimization with Gaussian processes
BSR	Buffer status reporting
BTS	Base transceiver station
CE	Control elements
CLPC	Closed loop power control
CPU	Central processing unit
CQI	Channel quality indicator
CSI	Channel state information
CTDE	Centralized training with decentralized execution
DDQN	Double-deep Q-network
DTC	Decision-tree calculation
eMBB	Enhanced mobile broadband
FD	Frequency domain
FPS	Fast protocol stack
FR1	Frequency range 1
HARQ	Hybrid automatic repeat request
InF	Indoor factory
JCAS	Joint communications and sensing
L2	Layer two
LCID	Logical channel ID
MAC	Medium access control
MAPPO	Multi-agent proximal policy optimization
MARL	Multi-agent reinforcement learning
MCS	Modulation and coding scheme
MIMO	Multiple input multiple output

ML	Machine learning
mMTC	Massive machine-type communications
MNO	Mobile network operator
MRSS	Multi-RAT spectrum sharing
MU-MIMO	Multi-user MIMO
OFDMA	Orthogonal frequency-division multiple access
OLLA	Outer loop link adaptation
OLPC	Open loop power control
PDU	Protocol data units
PF	Proportional fair (scheduler)
PHR	Power headroom reporting
PHY	Physical layer
PPO	Proximal policy optimization
QoS	Quality of service
RA	Random access
RACH	Random access channel
RAN	Radio access network
RL	Reinforcement learning
SAC	Soft-actor critic
SDPC	Soft-dropping power control
SINR	Signal-to-interference-and-noise ratio
SOHO	Small office, home office
TA	Timing advance
TD	Time domain
TTI	Transmission time interval
UE	User equipment
URLLC	Ultra-reliable low latency communications

References

[1] J. Hoydis, F. A. Aoudia, A. Valcarce, and H. Viswanathan, “Toward a 6G AI-native air interface,” IEEE Communications Magazine, vol. 59, no. 5, pp. 76–81, 2021.

[2] Nokia, “Taking 5G-Advanced to the next level and bridging into the 6G era,” 2023. [Online]. Available: https://onestore.nokia.com/asset/213705.

[3] Nokia, “Scaling up AI/ML for cellular radio access,” Oct 2023. [Online]. Available : https://www.nokia.com/blog/scaling-up-aiml-for-cellular-radio-access/.

[4] P. Kela and T. Veijalainen, “Cooperative action branching deep reinforcement learning for uplink power control,” in Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit), 2023.

[5] A. Brighente, J. Mohammadi, P. Baracca, S. Mandelli, and S. Tomasin, “Interference prediction for low-complexity link adaptation in beyond 5G ultra-reliable low-latency communications,” IEEE Transactions on Wireless Communications, vol. 21, no. 10, pp. 8403–8415, 2022.

[6] J. V. Evangelista, Z. Sattar, G. Kaddoum, B. Selim, and A. Sarraf, “Intelligent Link Adaptation for Grant-Free Access Cellular Networks: A Distributed Deep Reinforcement Learning Approach,” arXiv preprint arXiv:2107.04145, 2021.

[7] F. Geiser, D. Wessel, M. Hummert, A. Weber, D. Wubben, A. Dekorsy,¨ and A. Viseras, “DRLLA: Deep Reinforcement Learning for Link Adaptation,” in Telecom, vol. 3, no. 4. Multidisciplinary Digital Publishing Institute, 2022, pp. 692–705.

[8] S. Mandelli, A. Weber, P. Baracca, and J. Mohammadi, “TROLL: Training of Outer Loop Link Adaptation in Wireless Networks via Backpropagation,” in WSA 2021; 25th International ITG Workshop on Smart Antennas. VDE, 2021, pp. 1–6.

[9] P. Kela, T. Höhne, T. Veijalainen and H. Abdulrahman, “Reinforcement Learning for Delay Sensitive Uplink Outer-Loop Link Adaptation,” 2022 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit), Grenoble, France, 2022,

[10] L. Maggi, A. Valcarce, and J. Hoydis, “Bayesian optimization for radio resource management: Open loop power control,” IEEE Journal on Selected Areas in Communications, vol. 39, no. 7, pp. 1858–1871, 2021.

[11] A. Lieto, Q. Liao, and C. Bauer, “A generative approach for production aware industrial network traffic modeling,” in 2022 IEEE Globecom Workshops (GC Wkshps), 2022, pp. 575–580.

[12] Nokia, “”A novel approach to radio protocols design for 6G”. 2023 Retrieved from https://www.bell-labs.com/institute/white-papers/a-novel-approach-to-radio-protocols-design-for-6g/

[13] A. Valcarce and J. Hoydis, “Toward joint learning of optimal MAC signaling and wireless channel access,” IEEE Transactions on Cognitive Communications and Networking, vol. 7, no. 4, pp. 1233–1243, 2021.

[14] M. P. Mota, A. Valcarce, J.-M. Gorce, and J. Hoydis, “The emergence of wireless MAC protocols with multi-agent reinforcement learning,” in 2021 IEEE Globecom Workshops (GC Wkshps), 2021, pp. 1–6.

[15] L. Miuccio, S. Riolo, S. Samarakoon, D. Panno, and M. Bennis, “Learning generalized wireless MAC communication protocols via abstraction,” in GLOBECOM 2022 - 2022 IEEE Global Communications Conference, 2022, pp. 2322–2327.

[16] S. Seo, J. Park, S.-W. Ko, J. Choi, M. Bennis, and S.-L. Kim, “Towards semantic communication protocols: A probabilistic logic perspective,” IEEE Journal on Selected Areas in Communications, pp. 1–1, 2023.

[17] CENTRIC: Towards an AI-Native User-Centric Air Interface for 6G Networks. Retrieved from https://centric-sns.eu/

[18] Nokia, “Evaluation of ML for CSI feedback enhancement,” 3rd Generation Partnership Project (3GPP), TSG RAN WG1 contribution R1-2304681, 05.

[19] “Nokia Bell-Labs,” 2023.07.10 [online]. Simplifying spectrum migration from 5G to 6G. Available: https://www.bell-labs.com/institute/white-papers/simplifying-spectrum-migration-from-5g-to-6g/

[20] O. Naparstek and K. Cohen, “Deep multi-user reinforcement learning for distributed dynamic spectrum access,” IEEE Transactions on Wireless Communications, vol. 18, no. 1, pp. 310–323, 2019.

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

Document code: CID213703 (April)