

Toward a 6G AI-Native Air Interface

White paper

Each generation of cellular communication systems is marked by a defining disruptive technology of its time, such as orthogonal frequency division multiplexing (OFDM) for 4G or massive multiple-input multiple-output (MIMO) for 5G. Since artificial intelligence (AI) is the defining technology of our time, it is natural to ask what role it could play for 6G. While it is clear that 6G must cater to the needs of large distributed learning systems, it is less certain if AI will play a defining role in the design of 6G itself. The goal of this article is to paint a vision of a new air interface that is partially designed by AI to enable optimized communication schemes for any hardware, radio environment, and application.

Contents

Introduction	3
AI-Native Air Interface	4
Possible benefits	4
Three steps toward the AI-AI	6
Case study: From neural receivers to pilotless transmissions	7
The next frontier: protocol learning for the MAC	9
Learning a given protocol	10
Emerging a new protocol	11
Conclusion	12
Abbreviations	12
Authors	13
Sources	13
Acknowledgments	13
References	13

Introduction

While 5G is rolled out globally and the standardization discussions for its future evolution take place, researchers in academia and industry start to think about visions, use cases, and disruptive key technologies for a possible 6G system. Publicly funded 6G research projects in Europe [1], the United States and China are under way, and also the ITU has begun their work on requirements for fixed networks in the 2030s [2]. A common theme in many 6G vision papers is that of creating digital twin worlds for seamlessly connecting and controlling physical and biological entities to enable new mixed-reality super-physical experiences [3].

Apart from new spectrum technologies and the support of simultaneous communications and sensing as well as extreme connectivity requirements (among others), it is expected that machine learning (ML) and AI will play a defining role in the development of 6G networks end-to-end across the design, deployment and operational phases [4]. As the network evolves to programmable and flexible cloud native implementation, ML/AI-based network automation will be crucial to simplify network management and optimization. Networks will become “cognitive” in the sense that various aspects such as virtualized network function placement, slicing, quality of service, mobility management, radio resource management, and spectrum sharing will all rely on ML/AI to varying degrees. In fact, we expect that ML/AI will significantly impact even the 6G air interface, which is the focus of this paper.

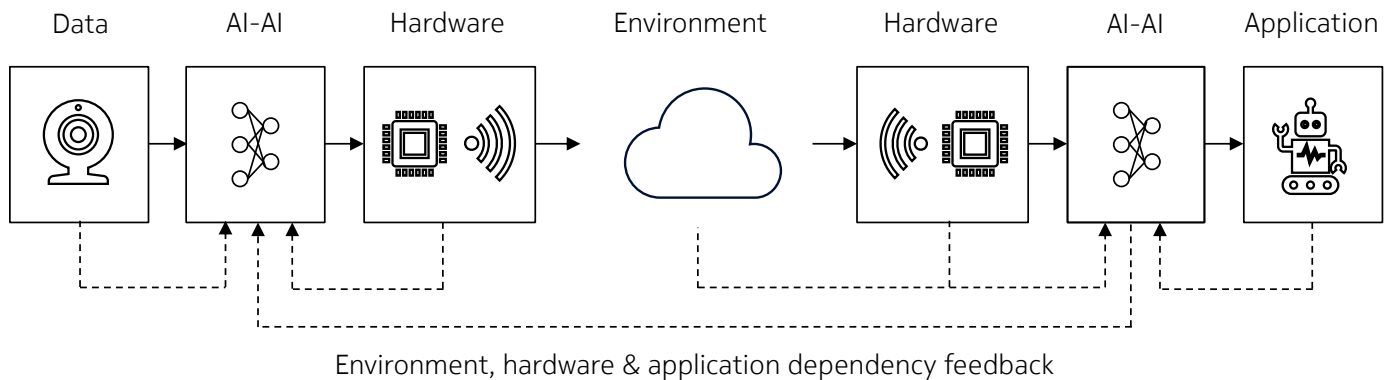
While ML starts to be widely used in the industry to enhance the implementation of various components within the 5G radio access network (RAN) and core, it is fair to say that there is not a single component of 5G that has been designed by ML. The purpose of this article is therefore to raise and discuss the question: What if 6G was designed in a way that ML/AI could modify parts of the physical (PHY) and medium access control (MAC) layers?

AI-Native Air Interface

Let us start by presenting our vision and motivation of an AI-Native Air Interface (AI-AI), which describes a disruptive change to the traditional way communication systems are designed, standardized and productized. We will first provide a summary of possible benefits of such an approach, then detail three important but also necessary development steps toward realizing our vision, and finally present a case study that exemplifies the respective potential performance gains and advantages.

As illustrated in Figure 1, the goal of the AI-AI is to serve an application with the data it needs in the most efficient way by taking into account the constraints of the available hardware and the radio environment. The AI-AI hence no longer decouples source and channel coding as well as communication of data from the intended use by an application, and embraces hardware constraints and undesired effects of the communications channel rather than fighting them. While the last decades were used to implement the scientific breakthroughs by Shannon and Wiener (as well as many others), we are now entering a new era for communications where classical approaches must be revisited and new theories developed to achieve the technological breakthroughs needed for a possible 6G system. We believe that our vision of an AI-AI could become a reality within the next decade by optimizing the air interface from end-to-end thanks to advances in the field of AI for communications.

Figure 1. The AI-Native Air Interface (AI-AI) adapts to different radio environments, hardware, data, and applications. Compared to previous air interfaces, it is not only designed to reliably transmit bits, but also to serve an application with the data it needs in an optimal way.



Possible benefits

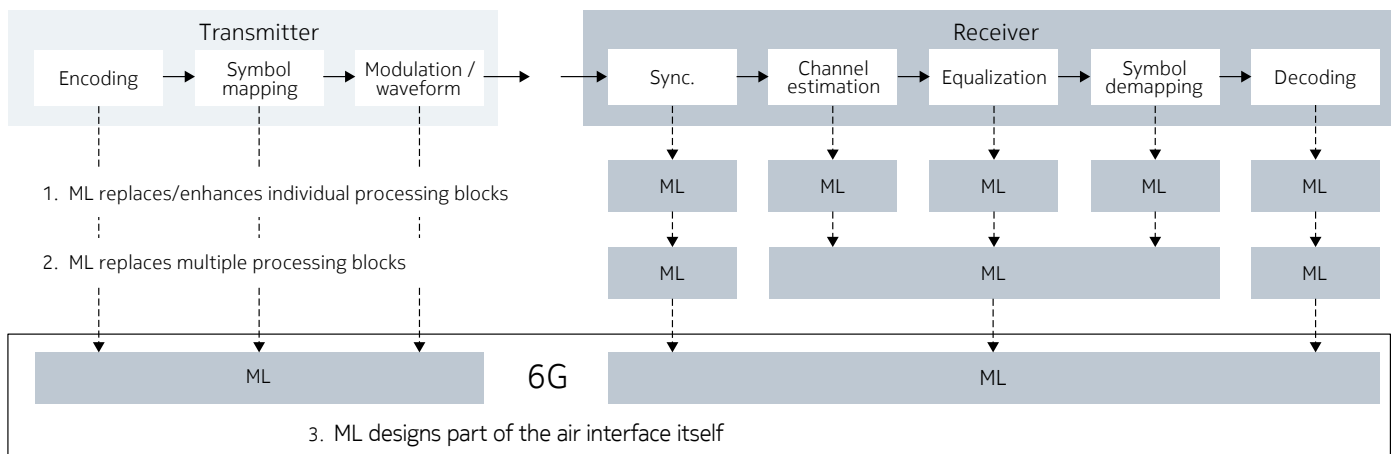
First, in contrast to a single classical waveform choice such as OFDM in 5G, the AI-AI could enable learning of bespoke waveforms for different frequencies, which not only make more efficient use of the spectrum but are also optimally adapted to practical limitations of the transceiver hardware and channel, such as non-linear power amplifiers, hybrid analog-digital processing, low quantization resolution, very short channel coherence time and bandwidth, phase and impulsive noise. Also new modulation schemes, pilot sequences, and codes can be learned or optimized with ML to squeeze even more performance out of the spectrum. Several research groups have recently demonstrated practical gains of such approaches (see for example, [5] and [6]). A very interesting application arises for the transmission of short messages, where the classical frame structure of preamble, pilots, coded bits, and cyclic-redundancy check could be replaced with a fully learned radio burst conveying a few bits of information. Also learning of new waveforms for simultaneous communication + “X” (e.g., sensing or power transfer) has high potential.

Second, fully learned transceivers have the benefit that they do not need to undergo the very costly and time-consuming traditional process of algorithm design and hardware implementation anymore. They can be trained directly for the targeted hardware platform (or even on it, depending on the capabilities). With an increasingly rich diversity of expected 6G use-cases and the emergence of small-scale sub-networks [7], this versatility becomes essential to ensure that 6G can cater for each individual use case and deployment scenario in the best way possible. Moreover, given the rapid speed at which ML hardware accelerators are developing, it is likely that learned transceiver implementations will rapidly outperform their traditional counterparts in power efficiency, latency, and cost. Advances in neuromorphic computing [8] could further amplify this trend.

Third, the more we follow the AI-AI principle of learning-based design and specification, the less needs to be standardized. The current 5G specification boasts a very rich sets of options and parameters for different frequency bands and scenarios, which pose a difficult challenge from an implementation point of view. It is undesirable to scale this approach to even more complex and diversified settings in 6G. If, on the other hand, only a sufficiently flexible framework for air interface learning was standardized, the system could auto-adjust to any kind of scenario. With a bit of wishful thinking, one could hope that 6G could be the last communication system to be standardized.

Fourth, the AI-AI allows integration of the data and the application consuming it into a single end-to-end learning process. Using the terminology from Shannon and Weaver’s seminal book [9], the AI-AI no longer only solves the problem of reliably transmitting bits (Level A), but simultaneously addresses the problems of semantics (Level B) and effectiveness (Level C) of communication (see, e.g.,[10]). While the latter aspects may not be applicable to the generic internet communication scenario, they become relevant for communication systems that are tailored to specific purposes and under the control of a single entity, such as industrial communication systems for sensing, surveillance and robot control.

Figure 2. Three phases toward the AI-AI: After gradually replacing most of the processing blocks, ML will design parts of the PHY layer.



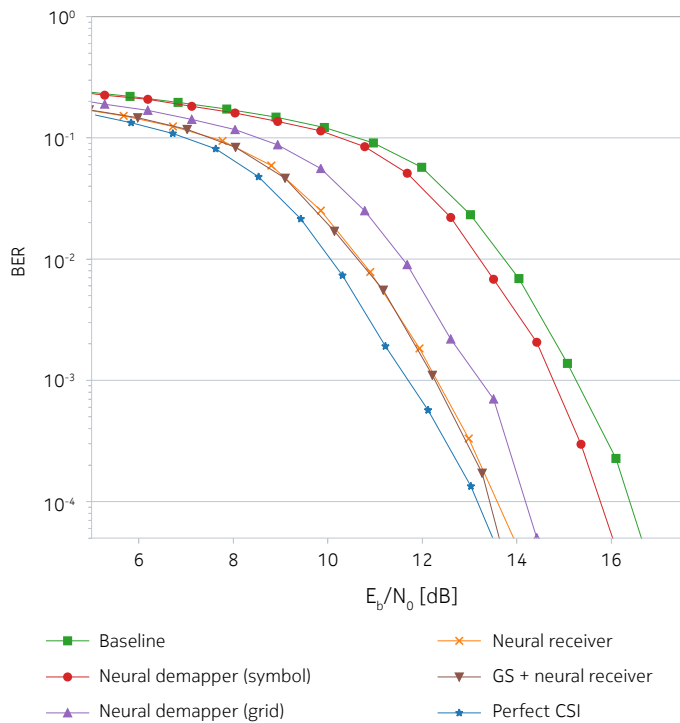
Lastly, the idea of end-to-end learning naturally extends to the MAC layer, where it would be desirable to emerge optimized signaling schemes and channel access policies that fluently transition from contention- to schedule-based depending on the use case and environment. Protocol learning could also address the problem of optimally multiplexing resources for communication and sensing (or other applications that radio waves can be used for). Ultimately, PHY and MAC layers could be jointly learned together.

Three steps toward the AI-AI

We currently see three important phases in the development and transition to the AI-AI, each of which requires sustained multi-disciplinary research. These are schematically shown in Figure 2. The first two phases do not require any new signaling or procedures as they only impact the implementation of transceiver algorithms. They can therefore be carried out on future 5G systems to gather practical experience while the research on 6G is progressing.

- 1) ML replaces single processing blocks: In the first phase, which is already happening in the industry today, ML will be used to enhance or replace some of the processing blocks, mostly in the receiver. Examples are physical random access channel detection, channel estimation, or symbol demapping. Although seemingly simple, this step constitutes a paradigm change in the way the industry designs and deploys radio transceivers. Even if the ML models are likely to be rather small, several important problems such as data acquisition, model updates and online training need to be solved, and hardware accelerators must be integrated into the PHY processing flow. The receiver processing will contain a mix of ML and traditional blocks.
- 2) ML replaces multiple processing blocks: In the second phase, more functionality is given to ML models, which take on the joint role of multiple processing blocks. This could be, for example, joint channel estimation, equalization and demapping. In this phase of the transition, the ML models will grow larger, hardware acceleration becomes increasingly important, and vendors need to commit to an “ML-only/ML-first” approach, because it is not viable to implement ML and non-ML backup solutions in parallel in the same processing platform due to increased power consumption and cost. This means that ML is also trusted more, although the inner workings of large models are less interpretable, but the potential gains are also higher. An example will be provided in Section II-C. In this phase, we will also realize what possibilities such learned transceiver components open up, for example, the need for fewer pilots, no cyclic prefix, and less stringent synchronization. In other words, we will learn what are the things ML allows us to do that we could not do before (with reasonable effort).
- 3) ML designs parts of the air interface: In the third phase, we will give even more freedom to ML/AI and let it design parts of the physical and MAC layers itself. This represents another paradigm change in the way communication systems are designed, because not all aspects of the PHY and MAC layers might be fixed in advance. This approach requires new forms of signaling and procedures to enable distributed end-to-end training. Rather than specifying modulation schemes and waveforms, for example, one would need to specify procedures that can be used to optimize these aspects of the air interface at deployment time. This is clearly something that nobody has ever done before and that requires a massive change in the way communication systems are standardized. It is of course also possible that ML-designed solutions to specific problems will be specified, which is already the case in 5G for the channel code design, for example.

Figure. 3. BER performance of all compared schemes in the case study



Case study: From neural receivers to pilotless transmissions


Next, we will present a case study that will go through the three phases toward the AI-AI outlined above and demonstrates the respective performance gains. We consider a doubly selective single-input single-output (SISO) channel at a carrier frequency of 3.5 GHz with the TDL-A power delay profile and a delay spread of 100 ns. The receiver is assumed to move at a speed of 50 km h⁻¹ and the channel evolves in time according to Jakes’ model. We consider cyclic prefix-based OFDM with 72 subcarriers spaced 30 kHz apart and assume transmission time intervals (TTIs) of 14 consecutive OFDM symbols, which contain codewords of length 1024 bits at a code rate of 2/3, generated by a 5G-compliant code. Our non-ML baseline assumes 64-quadrature amplitude modulation (QAM), pilots transmitted on every other sub-carrier on the third and twelfth OFDM symbols, least-squares channel estimation, equalization based on the nearest pilot, exact demapping to log-likelihood ratios (LLRs) assuming a Gaussian post-equalized channel, as well as a standard belief propagation (BP) decoder.

The BER performance of the baseline and all other schemes that will be subsequently introduced is shown in Figure 3. One can see that there is approximately a 3 dB gap between the baseline and a receiver, assuming perfect channel state information (CSI). We now describe some ways to close this gap using ML-enhanced receiver processing before delving into the benefits of optimizing parts of the transmitter too.

Due to channel aging and imperfect channel estimation, the quality of the post-equalized symbols that are fed into the demapper changes over the grid of resource elements (REs) within a TTI. A first possibility to cope with this problem is to learn a bespoke neural demapper for each RE (Phase 1). The BER performance of such a scheme is shown by the red line with dot markers in Figure 3. As expected, it provides some 0.5 dB gain over the baseline by computing better LLRs, but cannot compensate for channel aging, which results in a rotation and scaling of the equalized constellation.

In order to address these shortcomings, one can use a larger neural demapper which does not operate symbol-by-symbol but rather produces LLRs for the full TTI. It was shown in [11] and [12] that a fully convolutional ResNet architecture with dilated separable convolutions achieves remarkable performance for this task (see Figure 4). By having access to the full TTI of post-equalized symbols, the learned demapper can compensate for some of the errors made by the channel estimator and equalizer to provide a 2 dB improvement over the baseline (see purple line with triangular markers in Figure 3).

Figure 4. The neural receiver produces LLRs for an entire TTI of post-FFT symbols. Key architectural components are depth-wise separable dilated convolutions and residual connections. The same architecture can also be used as a neural demapper, operating on a TTI of equalized symbols.




Interestingly, it turns out that one can assign the joint task of channel estimation, equalization, and demapping to a neural network with a similar architecture (Phase 2). It is fed with a TTI of post-FFT received signals from which it directly computes LLRs for all symbols. In addition to the gains of the learned demapper, this neural receiver is now able to carry-out data-aided channel estimation and detection, resulting in an additional 0.5 dB gain. By increasing the model complexity and the size of the input (more sub-carriers and OFDM symbols), the performance can be brought arbitrarily close to the perfect CSI performance [11].

Lastly, we would like to investigate the benefits of a learned constellation (i.e., geometric shaping (GS)) at the transmitter side (Phase 3), which is jointly optimized together with the neural receiver. Figure 5 shows this constellation, which is used on every RE instead of the mix of pilots and 64-QAM symbols sent by the baseline. As can be seen from Figure 3, this system achieves the same BER as the neural receiver with 64-QAM, but has the additional benefit that no pilots are transmitted. End-to-end learning could hence remove the need and control overhead for demodulation reference signals.

This case study has only scratched the surface of what will be possible in the future. Interesting directions for future research include end-to-end learning for new waveforms, constrained hardware, very short messages, as well as joint source-channel coding for a specific application (which is also learned). Meta, transfer and federated learning are key enablers to make such schemes practical.

Figure 5. Learned constellation allowing pilotless transmissions together with a neural receiver. It has zero mean, unit power, and a single axis of symmetry. The optimal bit-labeling is also learned but not shown for readability.



The next frontier: protocol learning for the MAC

The protocols above the PHY take bit-by-bit transmission for granted to develop complex signaling schemes and orchestrate elaborate procedures across the network’s nodes. As a result, radios coordinate harmoniously and provide more powerful services than what point-to-point links offer.

The telecommunications industry defines these procedures through protocol standards, which are agreed upon in large meetings, where competing technical and economical interests are debated year after year. These efforts have a high cost and sometimes result in ambiguous technical specifications (TSs). After a TS is released, the implementation and test phase begins, which is even more costly. For this reason, it is interesting to question if this burdensome undertaking could be somehow automated. And if so, would the result perform better than human-designed protocols?

Wireless protocols are sequences of messages exchanged between radio nodes to transmit service data units (SDUs). As such, protocols can also be understood as a language between collaborative machines. Learning a language is something not only humans but also machines can do [13]. In fact, the field of learning to communicate (L2C) is growing fast thanks to recent developments in deep multiagent reinforcement learning (MARL) (see [14] for example). While most research efforts in this field are targeted toward natural languages, we believe these techniques can also be used for training wireless devices to learn communication protocols.

The 6G protocols of the AI-AI could be built this way and there are two ways to achieve this:

Learning a given protocol


Intelligent software agents can be trained to behave according to an a priori known protocol. Instead of coding the standard, agents are trained on it via machine learning. Such training would ideally be done only once during factory production. This yields a protocol implementation and, although it does not replace protocol standardization, it replaces protocol interpretation, implementation, and testing efforts. The cost savings and time-to-market reduction potential of 6G nodes built this way might be significant.

Already today, user equipment (UE) could be trained to learn the 5G New Radio (5G NR) MAC protocol. This would include learning to interpret the different control messages received from the base station (BS) (e.g., discontinuous reception (DRX), timing advance (TA)), as well as learning what to send in the uplink (e.g., buffer status report (BSR), power headroom report (PHR)). Learning can be formulated as a MARL problem, wherein the UE's MACs are deep reinforcement learning (RL) agents with two action spaces:

- Uplink signaling action space: All possible uplink control messages a UE may ever send.
- PHY action space: All channel access commands the MAC may invoke through the PHY application programming interface (API).

Protocol implementations trained this way may outperform expert systems thanks to the customized signaling and channel access policy [15]. The signaling is the vocabulary of messages the nodes have at their disposal as well as the rules about how to use them. The channel access policy decides how to make use of the PHY API, based on information it has thanks to the signaling. It is therefore fair to ask whether the gains are due to the learned signaling, the channel access policy, or both. The impact of signaling on the UE's actions can be quantified by metrics such as the instantaneous coordination (IC), which is the mutual information between the downlink signaling messages and the next channel access actions. This is illustrated in Figure 6, which shows that in unreliable channels, higher levels of coordination lead to performance gains (note the positive Pearson correlation coefficients ρ). The semantics of the MAC messages used in Figure 6 are fixed and known to the BS. Nevertheless, ML training can yield UE with widely different interpretations of these messages and consequently, also very different policies and performances (note the high variance of the results in Figure 6). This variance is a consequence of the vast size of the solution space, which strengthens the case for optimizing the protocols our radios use.

Figure 6. Relationship between performance and signaling-based coordination. Each point depicts the mean SDU rate and the IC of different learned UE MACs. Two UEs were trained to transmit two randomly arriving SDUs to the BS without collision or channel loss, via a protocol known to the BS and unknown to the UEs (i.e., UEs had no mapping between the protocol’s messages and PHY actions).



Emerging a new protocol

If there are that many possible ways a UE can communicate with a BS that implements a given protocol, why should we constrain ourselves to human-designed protocols? The next frontier in protocol learning is to let UE and BSs explore the entire space of possible protocols. This is challenging because for two radios to coordinate, they first need to find a state where they can interpret each other’s messages. Recent L2C research suggests that this language discovery problem may be overcome by first training the nodes with supervised learning. This essentially endows radios with an initial protocol that they can later evolve through self-play.

Communication protocols that emerge this way may be hard to interpret, which is essential for fault detection or performance monitoring. For this reason, some use cases may favor protocols that are close to known ones. This requires metrics that measure the distance between two protocols. Training techniques minimizing this distance may improve intelligibility.

The ability to learn new communication protocols opens the door to radio systems that are highly tailored to their deployment environment, thus boosting 6G capacities for niche and vertical markets. The AI-AI will not only reduce today’s signaling overheads, but also the standardization and development efforts for the highly complex radio technologies of the next decades. We foresee a 6G future where parts of the radio stack development cycle could be replaced by the click of a button.

Conclusion

While the next decade will tell if our vision of an AI-Native Air Interface provides sufficiently compelling benefits to make it into 6G, we are certain that AI/ML will profoundly change the way communication systems will be designed and deployed in the future. We hope that some of the readers will join us on this exciting journey.

Abbreviations

AI	artificial intelligence
AI-AI	AI-Native Air Interface
API	application programming interface
BER	bit error rate
BP	belief propagation
BS	base station
BSR	buffer status report
CSI	channel state information
DRX	discontinuous reception
FFT	fast Fourier transform
IC	instantaneous coordination
L2C	learning to communicate
LLR	log-likelihood ratio
MAC	medium access control
MARL	multiagent reinforcement learning
MIMO	multiple-input multiple-output
ML	machine learning
NR	New Radio
OFDM	orthogonal frequency division multiplexing
PHR	power headroom report
PHY	physical
QAM	quadrature amplitude modulation
RAN	radio access network
RE	resource element
RL	reinforcement learning
SDU	service data unit
SISO	single-input single-output
TA	timing advance
TS	technical specification
TTI	transmission time interval
UE	user equipment

Authors

Jakob Hoydis, Fayçal Ait Aoudia, Alvaro Valcarce, Harish Viswanathan

Sources

The content of the white paper mirrors a paper submitted to IEEE “Toward a 6G AI-native air interface”.

Acknowledgments

We would like to thank P. Mogensen, S. ten Brink, S. Cammerer, S. Dörner, M. Goutay, P. Srinath, M. Honkala, D. Korpi, J. Huttunen (and many others) for numerous discussions that helped to shape the vision outlined in this article.

References

- [1] 6G Waves 1, Spring 2020. 6G Flagship, University of Oulu, Finland, Mar. 2020.
- [2] ITU Focus Group on Technologies for Network 2030 (FG-NET- 2030), “Network 2030: A blueprint of technology, applications and market drivers towards the year 2030 and beyond,” White paper, May 2019.
- [3] H. Viswanathan and P. E. Mogensen, “Communications in the 6G era,” *IEEE Access*, vol. 8, pp. 57 063–57 074, Mar. 2020.
- [4] K. B. Letaief, W. Chen, Y. Shi, J. Zhang, and Y.-J. A. Zhang, “The roadmap to 6G: AI empowered wireless networks,” *IEEE Commun. Mag.*, vol. 57, no. 8, pp. 84–90, Aug. 2019.
- [5] S. Cammerer, F. Ait Aoudia, S. Dörner, M. Stark, J. Hoydis, and S. T. Brink, “Trainable communication systems: Concepts and prototype,” *IEEE Trans. Commun.*, Jun. 2020.
- [6] J. Downey, B. Hilburn, T. O’Shea, and N. West, “Machine learning remakes radio,” *IEEE Spectr.*, vol. 57, no. 5, pp. 35–39, Apr. 2020.
- [7] V. Ziegler, H. Viswanathan, H. Flinck, M. Hoffmann, V. Raïssaïnen, and K. Hattinen, “6G architecture to connect the worlds,” *IEEE Access*, vol. 8, pp. 173 508–173 520, Sep. 2020.
- [8] B. Rajendran, A. Sebastian, M. Schmuker, N. Srinivasa, and E. Eleftheriou, “Low-power neuromorphic hardware for signal processing applications: A review of architectural and system-level design approaches,” *IEEE Signal Process. Mag.*, vol. 36, no. 6, pp. 97–110, Nov. 2019.
- [9] C. E. Shannon and W. Weaver, *The mathematical theory of communication*. University of Illinois Press, 1949.
- [10] P. Popovski, O. Simeone, F. Boccardi, D. Gunduz, and O. Sahin, “Semantic-effectiveness filtering and control for post-5G wireless connectivity,” *J. Indian Inst. Sci.*, vol. 100, no. 2, pp. 435–443, May 2020.
- [11] M. Honkala, D. Korpi, and J. M. Huttunen, “DeepRx: Fully convolutional deep learning receiver,” arXiv:2005.01494, May 2020.
- [12] F. Ait Aoudia and J. Hoydis, “End-to-end Learning for OFDM: From neural receivers to pilotless communication,” arXiv:2009.05261, Sep. 2020.

- [13] T. B. et al., “Language models are few-shot learners,” in Proc. Int. Conf. Neural Inf. Proc. Sys. (NIPS), Vancouver, Spain, Dec. 2020.
- [14] J. N. Foerster, Y. M. Assael, N. de Freitas, and S. Whiteson, “Learning to communicate with deep multi-agent reinforcement learning,” in Proc. Int. Conf. Neural Inf. Proc. Sys. (NIPS), Barcelona, Spain, Dec. 2016, pp. 2145–2153.
- [15] A. Valcarce and J. Hoydis, “Towards joint learning of optimal signaling and wireless channel access,” arXiv:2007.09948, Jul. 2020.

About Nokia

We create the critical networks and technologies to bring together the world's intelligence, across businesses, cities, supply chains and societies.

With our commitment to innovation and technology leadership, driven by the award-winning Nokia Bell Labs, we deliver networks at the limits of science across mobile, infrastructure, cloud, and enabling technologies.

Adhering to the highest standards of integrity and security, we help build the capabilities we need for a more productive, sustainable and inclusive world.

For our latest updates, please visit us online www.nokia.com and follow us on Twitter [@nokia](https://twitter.com/nokia)

Nokia is a registered trademark of Nokia Corporation. Other product and company names mentioned herein may be trademarks or trade names of their respective owners.

© 2021 Nokia

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

Document code: CID210299 (February)