

6G Evolution Towards Connected Intelligent Machines

Dejan Vukobratovic¹ 

¹ Faculty of Technical Sciences, University of Novi Sad, Novi Sad, Serbia

dejanv@uns.ac.rs

Abstract – The evolution towards 6G raises a fundamental question: will future mobile networks connect machines more efficiently, or will they become an integral part of intelligent systems? This overview paper explores this question through the concept of task-oriented communications, which shifts the focus from faithfully transmitting data towards delivering the information needed to accomplish a task. We illustrate these ideas through two connected-robotics scenarios investigated in our research: semantic communication of 3D point clouds for industrial object recognition, and edge-based robotic localisation and mapping supported by 5G/6G connectivity. These examples highlight how communication, computation and sensing can be designed together to support potentially the fundamental role of 6G technology, that is to connect intelligent machines.

Keywords: 5G/6G, Task-oriented communications, Semantic communications, Connected robotics

1. INTRODUCTION

Mobile networks have evolved from connecting people through voice and data to supporting an increasingly diverse range of connected machines. As artificial intelligence (AI) advances, the next step materialised in 6G network standards may be to connect intelligent machines such as robots, autonomous vehicles and AI-driven systems that perceive, reason and act [1], [2]. This evolution raises new questions about what future networks should communicate, how they should process information, and how they can support the tasks performed by connected intelligent machines [3].

This paper explores these questions following the evolution of the emerging concept of task-oriented communications and focusing as use case examples on two connected-robotics scenarios investigated in our research. We first consider how the vision of 6G is evolving, then introduce the idea of communicating task-relevant information rather than simply reconstructing transmitted data. Finally, we illustrate this approach through industrial robotics communication and mobile-robot localisation using edge computing and 5G/6G infrastructure. Together, these examples offer a perspective on how communication, computation and sensing may become increasingly integrated as networks evolve towards 6G.

2. MOBILE NETWORK EVOLUTION TOWARDS 6G

Mobile networks have evolved from systems designed primarily for voice communication into platforms that connect people, machines, and services across the world. With 6G on the horizon, the question is not simply how to make these networks faster or more reliable, but whether they will need to play a different role altogether. As AI models increasingly move from data centres into robots, autonomous vehicles, and software agents, the network may become part of a much larger intelligent system rather than merely a means of transporting data. This section explores what that shift could mean, first by considering the emergence of intelligent machines and then by looking at how the vision of 6G is taking shape.

2.1 Connecting Intelligent Machines

In the past five decades, mobile cellular networks evolved through five generations of mobile systems. We are now at the brink of the sixth generation (6G). If we look back at the previous generations, it is not so easy to classify them. We often hear a common wisdom: the even generations were successful, while the odd generations failed. But perhaps this is looking at the history from a simplified perspective.

Let us recapitulate this evolution. 1G gave us mobile voice service. 2G made voice digital, scalable and truly global. 3G brought mobile data. 4G turned that data capability into seamless broadband connectivity. 5G shifted the ambition beyond connecting humans towards connecting machines. So perhaps the “odd generations” were not failures at all. Perhaps they were the generations that introduced new directions, while the following generation consolidated them and brought them to scale. If the above sequence of generations introduced voice - digital voice - data - broadband data - machines, what is the keyword 6G will bring in? If 6G is simply about making connected machines faster, more reliable, more efficient and more ubiquitous, this is exactly the expected consolidation and scaling step: communications has always progressed through a combination of technological innovation and massive engineering and economic scaling.

But if 5G is about connecting machines, the logical keyword for 6G could be connecting *intelligent machines*. Note that when 5G was conceived (early 2010s), the AI models and tools were only at their infancy and their evolution was, at the time, hard to predict. The distinction between connecting machines and connecting intelligent machines may be way more important than another increase in peak data rate.

We are entering a world in which machines are no longer just endpoints that generate data or execute commands. They increasingly perceive, reason, communicate and act. Some are physical, so called embodied intelligence, such as robots, autonomous vehicles, humanoids and other embodied intelligent systems [4]. Others are purely software, such as AI agents, assistants and autonomous digital services [5]. These two forms of intelligence may create very different demands on the network. Physical intelligent machines interact continuously with the physical world. Besides connectivity, they need sensing, perception, coordination and increasingly distributed intelligence. If these machines are no longer confined to a factory or a controlled environment, connectivity becomes considerably more complex problem, related not only to data rates and latency, but essentially to their capability to safely and reliably execute tasks. Software-based intelligent machines are different. They already run on our phones and in the cloud. Today their communication footprint may still be modest, but token exchange is growing very fast and that could change substantially depending on how human-AI and AI-AI interaction modes evolve.

2.2 Evolution of Mobile Network Standards Towards 6G


Crucial to the success of mobile networks lies in the global standardization coordinated by International Telecommunications Union Telecommunication Standardization Sector (ITU-T) and technically executed by 3rd Generation Partnership Project (3GPP). When 5G was being conceived in early 2010s, the key idea of what 5G would bring has been visually represented by famous “5G service triangle”. The idea was to extend 4G in three directions corresponding to three different services 5G network would offer. The principal service, called enhanced mobile broadband (eMBB), would extend the capabilities of 4G mobile broadband. The other two: massive machine-type communications (mMTC) and ultra-reliable low-latency communications (URLLC) were behind the 5G as the technology that aimed to connect the machines.

5G



Machines

6G



Intelligent Machines?

Figure 1. From 5G triangle to 6G hexagon [6]

In the past several years, conceptual design of 6G extended the notion of 5G triangle to the 6G hexagon (Figure 1). The idea is to extend 5G with three novel features: Integrated AI and Communication, Integrated Sensing and Communication, and Ubiquitous Connectivity. If 6G will be the technology to enable connected intelligent machines, then indeed, these three pillars constitute some of the key underlying features. But let us inspect this more closely.

With the aim of connecting intelligent machines, the main target of 6G is no longer “How much more capacity will 6G provide?” Instead, the main design problem can be broadly formulated as “What does the 6G network need to become when the endpoints themselves become intelligent?” This changes the role of the network.

Regarding the first new pillar of 6G: Integrated AI and Communication, in 3GPP, a lot of effort has been spent in discussing and standardizing the AI for the network, i.e., using machine learning and AI in optimising different network functionalities across the core network and the RAN such as scheduling radio resources, predicting traffic, automating operations and managing increasingly complex networks [7]. Significantly less effort is invested in standardizing the network for AI, although recently, frameworks such as AI/ML enablement (AIMLE) at the service level are arising to fill this gap [8]. As the network becomes integral part of the computing and intelligence infrastructure itself, communication, computation, sensing and AI increasingly interact rather than remain separate functions.

The second 6G pillar: Integrated Sensing and Communication (ISAC) is active topic of 3GPP standardization effort. We have moved from thinking of wireless systems purely as communication systems towards integrated sensing and communication. But for intelligent machines, we need to go one step further. The machine does not experience the world through RF sensing alone. It sees, hears, measures, localises and interacts through multiple sensing modalities. The network could become part of a much broader multimodal sensing infrastructure, with AI bringing these different observations together. Current 3GPP focus is leaned towards radar-like functionalities of future wireless interface, while most of the end-users (robotic, vehicular or similar industries) place focus on different modalities. This raises important question who will be the commercial user of 6G ISAC service in the future.

Finally, 6G targets ubiquitous connectivity. 5G already had the ambition of connecting essentially everything. Non-terrestrial networks (NTN) can extend that reach and help complete that vision. But for intelligent machines, ubiquitous connectivity may mean something more than simply having coverage everywhere. It means providing the connectivity, sensing and computational support wherever intelligent machines operate, including environments where infrastructure cannot be assumed. Thus it is reasonable to expect that the former two pillars will be of immediate importance for 6G connectivity of intelligent machines, while NTN and ubiquitous connectivity will complete the 5G ambition of connecting machines anytime and anywhere.

3. FROM COMMUNICATING DATA TO COMMUNICATING MEANING

Imagine a mobile robot navigating through a factory. Its cameras and LiDAR sensors continuously generate images and three-dimensional point clouds containing millions of measurements. Should the robot transmit all this information to a remote server at the edge or cloud, or could it transmit only the information needed to estimate its position, recognise an obstacle or plan its next movement? The answer depends on what we expect the communication system to accomplish.

This simple question captures the motivation behind task-oriented communications, also known as goal-oriented communications. Instead of treating the faithful error-free delivery of data as the ultimate objective, task-oriented communications starts from the task that the receiving system must perform. The communication system is then designed to preserve the information that matters for that task, potentially reducing the amount of data that needs to be transmitted while maintaining the required task performance.

3.1 From Bits to Tasks

Modern communications are built around a remarkably successful principle: represent information as bits, transmit those bits reliably, and reconstruct the original information exactly at the receiver. This approach follows the foundations of information and coding theory established by Claude Shannon [9]. Its success is visible everywhere, from mobile broadband and video streaming to industrial communications.

A conventional communication system typically separates the problem into several stages. First, source coding compresses the original data. Channel coding adds redundancy to protect the information against transmission errors. Modulation and waveform processing then prepare the information for transmission over a wireless channel. At the receiver, the corresponding operations recover the transmitted information.

This architecture works extremely well when the objective is to reproduce the original message, the goal established by Shannon. But intelligent machines introduce a different perspective. If a remote AI model needs to identify a particular object in a camera image, does it really need every detail of that image? If a robot needs to estimate its position, must the network deliver an exact reconstruction of every LiDAR measurement? Often, the answer is no.

What matters is not necessarily how accurately the original data can be reconstructed, but how well the receiving system can perform its task. A slightly distorted image may be perfectly adequate for recognising an object, while a visually excellent reconstruction might lose a small detail that is essential for robotic manipulation. The best representation of data depends on what we intend to do with it.

This is the central intuition behind task-oriented communications [3], [10]. Rather than optimising communication solely for bit-level accuracy or conventional measures of signal distortion, we consider the performance of the task itself. The relevant objective might be object-recognition accuracy, robot localisation error, the success rate of a manipulation operation or the stability of a control system. Communication system can be designed and its resources allocated according to how much they contribute to achieving that objective.

3.2 Learning What Information Matters

How can a communication system learn which information is important for a task? One possibility is to improve conventional compression. Images, video and three-dimensional point clouds can already be compressed before transmission, reducing the required bandwidth while attempting to preserve their important features. These techniques provide a useful starting point, but they generally aim to reconstruct the source data, rather than directly optimise the performance of a remote AI task.

A more radical approach is to design the communication and AI processing together. Deep learning makes this possible by representing different processing stages as interconnected neural networks. Instead of designing each stage independently, it becomes possible to train an end-to-end system to transform sensor measurements into a compact representation that can be transmitted over a wireless channel.

An important step in this direction is deep joint source-channel coding [11]. In contrast to a strictly separated compression and channel-coding architecture, a learned system can jointly adapt its

representation to the characteristics of the communication channel. Such approaches have demonstrated promising results, particularly when transmission conditions are challenging. However, if the objective remains the reconstruction of the original image or point cloud, the system is still fundamentally solving a data-reconstruction problem.

Task-oriented communications takes the idea further. Suppose a robot's camera captures an image containing a workpiece. Instead of transmitting the image for reconstruction, the system could extract features relevant to the downstream task and transmit those features to a remote AI model. Depending on the application, the transmitted representation might encode the object's position, its geometric characteristics or other information needed for recognition and manipulation. The receiver uses this representation directly, without necessarily reconstructing the original image.

The important distinction is that the transmitted representation is designed around the task, rather than around the appearance of the original data. It may therefore be compact, task-specific and very different from the sensor data itself. This approach usually referred to as semantic communications, which investigates how to represent and transmit information according to its meaning and intended use. In this context, semantics does not simply mean converting data into words or human-readable descriptions. For an intelligent machine, useful semantic information may consist of learned numerical features that preserve the information required by a particular AI model. However, the task-oriented communication is still an open problem, both from the perspective of theoretical limits, as well as the specific transmitter and receiver design (Figure 2).




Figure 2. Task-oriented communication system design challenges

3.3 Why This Matters for Connected Robotics

Connected robotics offers a natural setting for exploring these ideas. Robots increasingly rely on cameras, LiDAR, inertial sensors and other sources of information to perceive their surroundings, estimate their position and coordinate their actions. These sensors usually generate large volumes of data, while the processing required to interpret that data may be computationally demanding.

Wireless connectivity creates an opportunity to distribute or completely offload this processing between the robot, nearby edge computers and more powerful cloud infrastructure. However, sending all sensor data to a remote server can consume substantial communication resources and introduce delays. On the other hand, moving all computation onto the robot may be impractical when its size, cost, energy consumption or available processing power is limited.

Task-oriented communications offers a way to approach this trade-off. The robot can partly process its sensor data locally, extract information relevant to a particular task and transmit a compact representation to an edge AI service. The remote system can then perform the required inference or contribute to a larger processing pipeline. The appropriate division of work depends on the task, the

available computing resources, the wireless channel and the consequences of errors or delays [12]. This does not mean that transmitting raw data or reconstructing it at the receiver will become unnecessary. Some applications require detailed sensor information, and a representation optimised for one task may be unsuitable for another. A representation useful for object recognition, for example, may discard geometric details needed for precise manipulation. Task-oriented communication therefore introduces new design questions: which information should be retained, how should its representation be learned, and how should the system balance communication cost against task accuracy and reliability?

These questions remain active areas of research. There is no single semantic representation that is optimal for every task, nor is there yet a universally applicable mathematical theory and consequently design framework that determines the best transmitter, receiver and communication architecture for arbitrary tasks. An important challenge is to develop methods that preserve task-relevant information while remaining robust to changing wireless conditions and practical enough to integrate into real communication systems [12], [13].

As 6G networks increasingly support intelligent machines, communication can become more closely integrated with sensing, computation and control. Instead of treating sensor data as an end product that must be delivered intact, the network can become an active part of the process through which machines perceive their environment and accomplish their goals. In this sense, task-oriented communications represents a shift in perspective: from asking how faithfully we can transmit data to asking how effectively we can help a machine accomplish its task. The following section explores what this shift can mean in practice through connected-robotics use cases investigated jointly by the two teams at the University of Novi Sad - Faculty of Technical Sciences and the Institute of Artificial Intelligence Research and Development of Serbia.

4. CONNECTED ROBOTICS: FROM COMMUNICATING DATA TO UNDERSTANDING THE WORLD

The ideas behind task-oriented communication become easier to understand when we look at real robotic systems. A robot may need to recognise an object, estimate its position, or build a map of the space in which it moves. In each case, its sensors generate much more information than is ultimately needed for a particular decision. As noted in the previous section, sending all this information over a wireless network can consume valuable bandwidth and introduce delays. But what if we could design the communication system around what the robot or the edge AI server assisting it actually needs to know?

Connected Robotic Workcell

Hardware and Sensors

- Robotic hand (Fanuc CRX 10)
- Camera: 3D point cloud + stereo RGB image

Edge AI

- AI model: 3D object and pose reconstruction
- Motion planning & control



Connected Edge SLAM

Hardware and Sensors

- UGV platform
- Wheel odometry, lidar, camera, CSI data

Edge AI

- Edge model: Factor Graph SLAM running Gaussian BP



Figure 3. Two representative 6G connected robotics use cases

Our research explores this question through two specific connected-robotics scenarios. The first concerns the transmission of three-dimensional information about objects in industrial environments. The second investigates how mobile robots can efficiently perform edge-based simultaneous localization and mapping (SLAM) while distributing their perception and localization tasks between the robot itself and nearby edge servers, while exploring how the wireless network can contribute as an additional sensing modality.

4.1 Semantic Communication of Point Clouds

In modern industrial environments, robots increasingly rely on depth cameras and other three-dimensional sensors to understand objects and their surroundings. These sensors can produce a point cloud: a collection of points representing the surfaces of objects in three-dimensional space. Point clouds can describe a scene in considerable detail, but transmitting them over a wireless network is demanding, particularly when the information must arrive quickly enough to support an ongoing robotic task.

Traditional compression approaches attempt to reduce the amount of data while preserving the geometry of the original point cloud as accurately as possible. This makes sense when the objective is to reconstruct or visualise an object. But consider a different situation: a robot receives information from a remote sensor because it needs to identify a component or determine whether a particular object is present. Does it really need every surface point to be reconstructed precisely? Or is it more important that the information arriving at the other end supports the correct decision?

These questions motivate our work on task-oriented wireless transmission of 3D point clouds. We are working towards developing a learning-based communication approach in which the transmitter converts a point cloud into a compact representation, which is sent over a simulated wireless channel and processed at the receiver. The system is designed to support two different objectives: reconstructing the object's geometry and recognising what the object is. This allows us to study how these objectives respond differently to imperfections in wireless transmission.

The experiments use data from an industrial setting, including the recognition of objects relevant for autonomous welding task. The results reveal an important behaviour. As the wireless channel deteriorates, the reconstructed geometry can become noticeably less accurate, while object recognition remains comparatively stable over a broad range of channel conditions. In other words, small errors in the reconstructed positions of individual points do not necessarily prevent a learning model from recognising the object correctly [14].

This finding illustrates one of the central ideas of semantic communication: the quality of a communication system should be judged not only by how accurately it reproduces the transmitted data, but also by how well it supports the task for which that data is needed. In an industrial application, the ability to make a reliable decision may matter more than reproducing every geometric detail.

There is an important qualification, however. Geometry still matters when a robot needs precise object dimensions, accurate positioning, or reliable manipulation. The goal is not to discard geometric information indiscriminately, but to understand which information must be preserved for a given task. A representation suitable for recognising an object may not be sufficient for grasping it. The challenge is to find the right balance between communication cost, robustness, and the requirements of the application.

4.2 Edge SLAM for 6G Connected Robotics

A second research scenario concerns mobile robots that continuously determine where they are and what their surroundings look like. This capability is known as simultaneous localisation and mapping, or SLAM. It allows a robot to estimate its position while building a map of an unfamiliar environment, using observations from sensors such as cameras, lidars and wheel odometry.

SLAM can require substantial computing power, particularly when it involves processing detailed lidar and camera images or combining several sources of information. A limited-compute mobile robotic platform cannot always perform every computation locally while meeting its performance and energy constraints. One option is to transfer some of the processing to a nearby edge server connected through

a 5G/6G network.

This arrangement introduces a challenge. The robot should send its sensor data across the network, the edge server must process it, and the resulting information should arrive in time to be useful. The wireless link is therefore not simply a transparent pipe between two endpoints: its latency, data rate, and behaviour affect the overall robotic system. Even the software mechanisms, such as widely used ROS2 middleware, that is used to package and transport sensor messages can influence end-to-end performance.

To investigate these interactions, our team has developed a connected mobile-robot platform integrating an unmanned ground vehicle, a 5G Open RAN testbed, and an edge server running a SLAM application. The platform allows us to examine and keep full observability of the complete path from the robot's sensors, through the robotic software and 5G network, to the computation performed at the edge. This makes it possible to investigate the system as a whole rather than studying the wireless link and the robotics software independently.

One part of this research focuses on a particularly interesting question: where should the processing take place? For example, instead of transmitting an entire camera image for every subsequent operation to be performed at the edge server, a learning-based visual-processing pipeline can be split between the robot and the edge computer. The robot performs an initial part of the computation and sends an intermediate representation to the edge, where the remaining processing is completed.

Recently, we have explored this idea for the detection of visual markers used to estimate the robot's position [15]. The learning model can be split at different points, changing how much computation takes place onboard and how much information must be transmitted. The intermediate representation is not intended to recreate the original image; it carries features needed for the remaining visual-processing task. The choice of split point therefore creates a trade-off between local computation, communication overhead, and the accuracy of the final localisation result, while the overall design goal is to minimize end-to-end latency while keeping the task performance at the satisfactory level. The communication system is no longer designed solely around the transmission of camera images. Instead, the transmitted information is tied to the computation that must be completed at the other end.

The same research platform opens the door to an even broader idea. A wireless network may do more than transport information between a robot and an edge server. The radio signals used for communication also interact with the surrounding environment, and measurements of these signals can contain information about that environment, as part of integrated sensing and communication paradigm. In particular, the wireless system can provide channel state information (CSI), which describes characteristics of the radio channel between the network and a connected device. Such measurements may complement information from cameras, odometry, and other sensors.

Our current work investigates how radio-channel information can be made available to the robot's SLAM application through the 5G Open RAN infrastructure. This is not simply a matter of switching on an existing feature: the radio measurements must be extracted from the network, delivered to the robotic software, and associated correctly with the robot's other sensor observations. We have developed and tested a mechanism for exposing CSI from the radio processing system to the edge application, while exploring an alternative based on the network's programmable control infrastructure. Using these measurements to improve the SLAM algorithm itself remains a further research step.

The longer-term vision is a more tightly integrated system in which the robot, edge server, and wireless network cooperate in understanding the environment. Rather than treating communication, computation, and sensing as separate functions, the system could exploit their interaction to support more capable and adaptable robots.

5. CONCLUSIONS: CONNECTING INTELLIGENT MACHINES

Is 6G going to be evolutionary or revolutionary? Perhaps the more useful question is what we expect future mobile networks to make possible. If 5G expanded the role of mobile networks from connecting people to connecting machines, 6G may take the next step towards connecting intelligent machines that must continuously perceive, reason, coordinate, and act in the physical world.

In this paper, we presented two examples that illustrate two complementary directions in our research. In the first, we investigate how to transmit information that preserves the properties needed for an industrial perception task, even when the wireless channel is imperfect. In the second, we explore how

to distribute robotic processing across the robot and the edge, and how information generated by the wireless network itself might contribute to environmental perception.

Both directions point beyond the traditional view of a mobile network as a system that simply delivers bits from one endpoint to another. For connected robotics, communication has to be considered alongside the computation and sensing that give those bits their purpose. The most useful information is not necessarily the most detailed, and the best place to process it is not necessarily on the robot or in the cloud alone.

These are still active research questions, rather than problems with universal solutions. Different robotic tasks demand different levels of geometric accuracy, latency, and robustness; wireless conditions and computing resources also change. The opportunity lies in designing these elements together, so that communication supports what the robot is trying to achieve.

This perspective brings us back to the broader question raised by the evolution towards 6G. If future networks are to support increasingly intelligent machines, success may depend not only on how much data they can carry, but also on how effectively they help those machines perceive, decide, and act.

ACKNOWLEDGEMENTS

This paper is based in part on the author's keynote presentation at the Global 6G Conference in April 2026 in Nanjing, China, and in part on the author's presentation at the panel entitled "Is 6G going to be evolutionary or revolutionary?" at the IEEE CSCN 2026 in London, UK. The author would like to thank colleagues from One6G WG2, work item WI216 on 6G connected robotics, and colleagues from IEEE P1955 standardization initiative on 6G connected robotics for insightful discussions. Finally and most importantly, the author would like to thank to the collaborators from research teams at the University of Novi Sad, Faculty of Technical Sciences, and The Institute of Artificial Intelligence Research and Development, who contributed to the results described in this paper.

DECLARATIONS

Conflict of interest (mandatory): The authors declare no conflict of interest.

Data availability (mandatory): The data availability is not applicable.

Author contributions (mandatory): The sole author contributed to this paper.

Ethical approval (mandatory for research involving human participants, human data or animals): The ethical approval is not applicable.

REFERENCES

- [1] M. Ghassemian, A. M. Valenzuela, A. G. Armada, D. Vukobratovic, P. Chatzimisios, K. Althofer, P. R. Venkatesh: "6G Empowering Future Robotics: A Vision for Next-Generation Autonomous Systems," to appear, IEEE Communication Magazine, 2026.
- [2] G. Singh, S. Amatare, E. Natalizio, D. Roy: "Realizing the 6G-Enabled Industrial Internet of Robotic Things: Targets, Enablers, and a Case Study," to appear, IEEE Internet of Things Magazine, 2026.
- [3] D. Gündüz, Z. Qin, I.E. Aguerri, H.S. Dhillon, Z. Yang, A. Yener, K.K. Wong, C.B. Chae, "Beyond transmitting bits: Context, semantics, and task-oriented communications," IEEE Journal on Selected Areas in Communications, Vol. 41, No. 1, pp. 5-41, 2022.
- [4] L. Bariah, M. Debbah: "AI embodiment through 6G: Shaping the future of AGI," IEEE Wireless Communications, Vol. 31, No. 5, pp.174-181, 2024.
- [5] A. Antonopoulos, E. Kartsakli, N. Bartzoudis, D. Brodimas, D. Vukobratovic, D. Tsolkas, F. Diego, G. Roumelas, L. Tomaszewski: "AGILE-6G: Agentic AI for autonomous management of 6G network/application services," to appear, IEEE Network, 2026.
- [6] <https://www.itu.int/en/ITU-R/study-groups/rsg5/rwp5d/imt-2030/pages/default.aspx> (Accessed: 30.09.2026)
- [7] X. Lin: "An overview of the 3GPP study on artificial intelligence for 5G new radio," arXiv preprint arXiv:2308.05315, 2023.
- [8] P. Taksande, S. Kiran, P. Jha, P. Chaporkar: "AI/ML in 3GPP 5G Advanced: Services and Architecture," arXiv:2512.03728, 2026.
- [9] C.E. Shannon: "A mathematical theory of communication," The Bell system technical journal, Vol. 27, No. 3, pp.379-423, 1948.
- [10] Y. Shi, Y. Zhou, D. Wen, Y. Wu, C. Jiang, K.B. Letaief: "Task-oriented communications for 6G: Vision, principles, and technologies," IEEE Wireless Communications, Vol. 30, No. 3, pp.78-85, 2023.

- [11] E. Bourtsoulatze, D.B. Kurka, D. Gündüz: "Deep joint source-channel coding for wireless image transmission," *IEEE Transactions on Cognitive Communications and Networking*, Vol. 5, No. 3, pp. 567-579, 2023.
- [12] V. Ninkovic, D. Vukobratovic, D. Miskovic, M. Zennaro: "COMMSPLIT: A Communication-Aware Split Learning Design for Heterogeneous IoT Platforms," *IEEE Internet of Things Journal*, Vol. 12, No. 5, pp. 5305-5319, March 2025.
- [13] V. Ninkovic, D. Vukobratovic, D. Miskovic, C. Wang: "Adaptive Unequal Error Protection for Semantic Split Learning over Wireless Channels," to appear, *IEEE Communication Letters*, 2026.
- [14] V. Ninkovic, T. Sobot, V. Vincan, G. Gojic, D. Miskovic, D. Vukobratovic: "Task-Oriented Wireless Transmission of 3D Point Clouds: Geometric Versus Semantic Robustness," In *2026 Global 6G Conference (Global 6G)*, pp. 1-6, Nanjing, China, April 2026.
- [15] B. Radovanovic, V. Ninkovic, K. Vidojevic, B. Bajic Papuga, D. Vukobratovic: "A Semantic Communication Approach to Fiducial Marker Processing in 5G-Enabled Edge SLAM," In *IEEE CSCN 2026*, London, UK, September 2026, arXiv preprint arXiv:2608.09620.

AUTHOR BIOGRAPHIES



Dr Dejan Vukobratović received a PhD.degree in electrical engineering from University of Novi Sad, Serbia, in 2008, where he is currently a Full Professor. During 2009-2010, he was Marie Curie postdoctoral fellow at the University of Strathclyde, Glasgow, UK. He has published more than 60 journal papers and 120 conference papers in top-tier IEEE journals and conferences. He received a best paper award at IEEE MMSP 2010 and IEEE SmartGridComm 2017. He serves as an Associated Editor for *IEEE Transactions on Communications*. He was TPC Co-Chair of IEEE VTC Spring 2020 (Antwerp) and IEEE SmartGridComm 2022 (Singapore), Symposia Chair on IEEE SmartGridComm 2021 (Aachen) and IEEE ICC 2023 (Rome), and IEEE CSCN 2026 (London), and General Chair for BalkanCom 2021 (Novi Sad), IEEE CSCN 2024 and FCN 2025 (Belgrade). His research interests include wireless communications, signal processing, machine learning and information theory applied in mobile cellular systems (5G/6G) and IoT communications.