

APPLIED RESEARCH

RLOps: Development Life-Cycle of Reinforcement Learning Aided Open RAN

PEIZHENG LI¹, (Graduate Student Member, IEEE), JONATHAN THOMAS¹,
XIAOYANG WANG¹, AHMED KHALIL¹, ABDELRAHIM AHMAD², RUI INACIO²,
SHIPRA KAPOOR³, ARJUN PAREKH³, ANGELA DOUFEXI¹, (Member, IEEE),
ARMAN SHOJAEIFARD⁴, (Senior Member, IEEE), AND ROBERT J. PIECHOCKI¹

¹Department of Electrical and Electronic Engineering, University of Bristol, BS8 1TH Bristol, U.K.

²Vilicom U.K. Ltd., RG1 8HH Reading, U.K.

³Applied Research BT, IP5 3RE Suffolk, U.K.

⁴InterDigital Communications Inc., Wilmington, DE 19809, USA

Corresponding author: Peizheng Li (peizheng.li@bristol.ac.uk)

This work was supported in part by the Innovate UK/CELTIC-NEXT European collaborative project on AI-enabled Massive MIMO (AIMM), in part by the Next-Generation Converged Digital Infrastructure (NG-CDI) Project, and in part by the BT and Engineering and Physical Sciences Research Council (EPSRC) under Grant EP/R004935/1.

ABSTRACT Radio access network (RAN) technologies continue to evolve, with Open RAN gaining the most recent momentum. In the O-RAN specifications, the RAN intelligent controllers (RICs) are software-defined orchestration and automation functions for the intelligent management of RAN. This article introduces principles for machine learning (ML), in particular, reinforcement learning (RL) applications in the O-RAN stack. Furthermore, we review the state-of-the-art research in wireless networks and cast it onto the RAN framework and the hierarchy of the O-RAN architecture. We provide a taxonomy for the challenges faced by ML/RL models throughout the development life-cycle: from the system specification to production deployment (data acquisition, model design, testing and management, etc.). To address the challenges, we integrate a set of existing MLOps principles with unique characteristics when RL agents are considered. This paper discusses a systematic model development, testing and validation life-cycle, termed: RLOps. We discuss fundamental parts of RLOps, which include: model specification, development, production environment serving, operations monitoring and safety/security. Based on these principles, we propose the best practices for RLOps to achieve an automated and reproducible model development process. At last, a holistic data analytics platform rooted in the O-RAN deployment is designed and implemented, aiming to embrace and fulfil the aforementioned principles and best practices of RLOps.

INDEX TERMS O-RAN, machine learning, reinforcement learning, MLOps, RLOps, digital twins, data engineering.

I. INTRODUCTION

As the forefront of a mobile communication network, the Radio Access Network (RAN) directly interacts with the user equipment (UE). Its architecture has undergone profound changes in recent years, transitioning from monolithic to disaggregated architectures and from vendor-based to open-source solutions [1]. The disaggregation of the RAN is reflected in two vectors, one is the horizontal disaggregation

The associate editor coordinating the review of this manuscript and approving it for publication was Li He¹.

of the network functions with open interfaces, and the other is the virtualization of hardware and software in vertical. To achieve efficiency dividends, allow for increased innovation, and also performance gain, O-RAN¹ emerged from years of industry work in groups studying possible Open RAN trends (including the 3rd Generation Partnership Project (3GPP)). O-RAN is based on 3GPP new radio (NR)

¹In this paper, O-RAN is taken as the reference architecture to demonstrate the validity of the proposed RLOps, but the principles can be applied to others RAN architectures.

specifications, meaning it is 4G and 5G compliant, while the difference is merely the additional interfaces that are defined by O-RAN, focusing on a functional split called 7.2x.

The emphasis of O-RAN has been on *openness* and *intelligence* from the beginning [2], through which it intends to actively embrace the technological revolution brought by machine learning (ML). Within recent years significant research has been undertaken demonstrating the potential of ML within telecommunications, including channel estimation in massive multiple-input and multi-output (MIMO) systems [3], resource and service management in large-scale mobile ad-hoc networks [4] and mobile edge computation offloading and edge caching [5], to name but a few. The introduction of this class of methods is supported through a number of avenues but perhaps most importantly through the definition of open interfaces and radio intelligence controllers (RICs). Through their introduction, O-RAN provides the foundations by which ML models can be introduced into RAN. Thereby, facilitating the evolution of RANs from being static and stiff to data-driven, dynamically sensing and self-optimising. Notably, the exact mechanisms and procedures required to deploy cutting-edge ML solutions into production and realise their economic potential still need clarification.

As ML models and systems continue to mature they are experiencing increased adoption within a range of industrial settings. The process through which they are developed and deployed is being formalised under the banner of MLOps [6]. Whereby, this is comparable to DevOps [7] and emphasises similar best practices whilst considering the unique challenges which the relevance of data in ML model development introduces. In order to reliably and consistently bring the potential of ML to O-RAN, an operational platform implementing MLOps principles whilst considering the unique challenges of RANs is required. These challenges pertain to its prominence as the critical national infrastructure and the highly dynamic nature of the platform. We place particular emphasis on the challenges of Reinforcement Learning (RL) within O-RAN due to the numerous applications that exist for it and its relative immaturity in terms of industrial applications. Where we discuss key elements throughout the applications lifecycle including design considerations, challenges pertaining to training within the simulation and effectively monitoring live deployments, to name but a few. This pipeline and associated set of principles is coined *RLOps*.

In this paper, we introduce the principles and best practices of RLOps in the context of an O-RAN deployment [2]. To the best of our knowledge, this is the first work to systematically discuss the life-cycle development pipeline of ML, especially RL models in O-RAN and put forward a network analytics platform in accordance with RLOps. We explain the fundamental principles and highlight critical factors involved in RLOps. In Section II, we briefly introduce ML and RL, and the evolution of RAN and O-RAN architectures are given in detail. Next, we discuss some related applications of ML/RL in intelligent O-RAN, and correspondingly a series of challenges encountered in the development and deployment

stages of ML/RL models. In Section III, we elaborate on the principles of RLOps from the perspective of design, deployment and operations. We highlight the safety and security concerns in RLOps. In Section IV, we put forward the effective routines and best practices of operating the aforementioned principles from the view of digital twins, automation and reproducibility. Section V illustrates the O-RAN related data analytics platform designed for achieving the principles and best practices of RLOps. Finally, Section VI concludes this paper. Table 1 gives the list of used acronyms.

II. RELEVANT BACKGROUND

A. MACHINE LEARNING IN GENERAL

Machine Learning (ML) is a branch of Artificial Intelligence (AI) concerned with learning from data, e.g. supervised learning (SL) and unsupervised learning (UL), or interaction, e.g. RL [8]. In general, ML considers the utilization of an adaptive model parameterized by θ with the intention of minimizing some objective function $\mathcal{J}(\theta)$. The exact objective function and form of the adaptive model depend on the exact formulation of the task (or set of tasks) we are interested in.

Within SL tasks, we are typically presented with a dataset $\mathcal{D} = \{x_i : y_i\}_{i \in \mathcal{D}}$ comprising of feature vectors x_i , which are labelled y_i which may refer to a discrete categories (cat or dog, for example) in a classification task or a real number if it is a regression task. Within this class of problem the objective is to learn a mapping $\mathcal{F}_\theta : \mathcal{X} \rightarrow \mathcal{Y}$. A typical formulation of our objective function is the minimizing of Negative log-likelihood in the case of classification or the Mean Squared Error in the case of regression tasks.

Like SL, in UL we are typically presented with a dataset $\mathcal{D} = \{x_i\}_{i \in \mathcal{D}}$, but in this case there are no labels. When presented with a task of this nature, we may be interested in clustering [9], density estimation [10] or in dimensionality reduction for visualization [11].

RLs interaction with data is fundamentally different from other forms of ML. Typically, the problem is formalised as a Markov Decision Process (MDP), where this is defined by the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma \rangle$. Where \mathcal{S} is the set of environment states, \mathcal{A} is the set of actions that an agent performs, \mathcal{P} represents the transition probability from any state $s \in \mathcal{S}$ to any state $s' \in \mathcal{S}$ for any given action $a \in \mathcal{A}$. \mathcal{R} is the reward function that indicates the immediate reward received from the transition from s to s' , and γ is the discount factor that trades off the instantaneous and future rewards. The intention is to find a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ which maximises the expected cumulative discounted reward \mathcal{G} [12] as defined in Equation 1.

$$\mathcal{G} = \mathbb{E}[\sum_{t \geq 0} \gamma^t \mathcal{R}(s_t, a_t, s_{t+1} | a_t \sim \pi(\cdot | s_t), s_0)] \quad (1)$$

The process of finding this π requires exploratory behaviours such that the agent can evaluate policies and learn about the MDP. The parameterization of the adaptive model may vary; for example, in model-free algorithms, we may parameterise

TABLE 1. List of acronyms.

Acronym	Definition	Acronym	Definition
3GPP	3rd Generation Partnership Project	ML	Machine Learning
5G	5th Generation Mobile Network	mMTC	Massive Machine-type Communications
AI	Artificial Intelligence	NE	Network Elements
API	Application Programming Interface	Near-RT RIC	Near-real Time RAN Intelligent Controller
BBU	Baseband Unit	NFV	Network Function Virtualization
BS	Base Station	NLP	Natural Language Processing
CAPEX	Capital Expenditure	Non-RT RIC	Non-real Time RAN Intelligent Controller
CPO	Constrained Policy Optimization	OFDM	Orthogonal Frequency-division Multiplexing
CI/CD	Continuous Integration (CI) and Continuous Deployment	OPEX	Operating Expense
CMDP	Constrained MDP	O-RAN	Open RAN
C-RAN	Centralized (or Cloud) Radio Access Network	PL	Processing Layer
CSI	Channel State Information	PPO	Proximal Policy Optimization
CU	Centralized Unit	PRB	Physical Resource Block
CV	Computer Vision	QoE	Quality of Experience
DCA	Data Collection Agents	QoS	Quality of Service
DT	Digital Twins	RAN	Radio Access Network
DL	Deep Learning	RF	Radio Frequency
DML	Data Mediation Layer	RIC	RAN Intelligent Controller
D-RAN	Distributed RAN	RL	Reinforcement Learning
DQN	Deep Q Learning	RNN	Recurrent Neural Network
DRL	Deep Reinforcement Learning	RRHs	Radio Remote Heads
DSL	Data Storage Layer	RU	Radio Unit
DU	Distributed Unit	SARSA	State-action-reward-state-action
eMBB	Enhanced Mobile Broadband	SDN	Software-defined Networking
ETSI	European Telecommunications Standardization Institute	SL	Supervised Learning
FCAPS	Fault, Configuration, Accounting, Performance, Security	SLA	Service Level Agreement
FDD	Frequency Division Duplexing	SMO	Service Management and Orchestration
IID	Independent and Identically Distributed	SON	Self-Organizing Network
IP	Internet Protocol	SOTA	State-of-the-art
IRS	Intelligent Reflecting Surfaces	TTI	Transmission Time Interval
LSTM	Long Short Term Memory	UAV	Unmanned Aerial Vehicle
LDPC	Low-density Parity-check Code	UE	User Equipment
MCTS	Monte Carlo Tree Search	UL	Unsupervised Learning
MDP	Markov Decision Process	URLLC	Ultra-reliable low-latency Communications
MEC	Mobile Edge Computing	V2X	Vehicle-to-everything
MIMO	Multiple-input and Multiple-output	VNFs	Virtual Network Functions
		vRAN	Virtual Radio Access Network

our π directly or the state-action value Q , or in model-based algorithms, we may learn a model of the MDP directly.²

B. EVOLUTION OF THE RAN AND O-RAN ARCHITECTURE

A typical mobile communication network mainly comprises a RAN, a transport network and a core network. The RAN gives the UE access to the core network, this subsequently provides the services to the user. The transport network implements the IP routing and IPsec functionality that securely connect the different network elements and network domains of the mobile network, thus allowing for full end-to-end functionality.

From 1G to 5G, the evolutionary trends of communication systems are the modularity and virtualization of decoupled network functionalities. For instance, the core network embraces x86 platform universal servers and performs network function virtualization (NFV), where the slicing of the core network embodies this feature. However, due to the complexity of antennas, the Remote Radio Heads (RRHs), and the Baseband Units (BBUs) in RAN, the functionality decoupling of RAN is slower than the decoupling of the

transport network and core network. Three distinct structural improvements have been proposed in the evolution of RAN, namely the distributed RAN (D-RAN), the centralized (or cloud) RAN (C-RAN), and vRAN. The RRHs and BBUs are co-located in D-RAN at every distributed cell site. The RRHs and BBUs communication are provided by the proprietary interfaces. Cells are connected back to the core network through the backhaul interface. In C-RAN, all BBUs are further concentrated into the centralized BBU pool for cloudification, and every site merely keeps antennas and RRH. RRHs and the centralized BBU are connected with fronthaul. Centralized BBUs bring the convenience of cell deployment and maintenance and significantly reduce the CAPEX and OPEX. vRAN decouples the software and hardware by NFV, where the BBU is virtualized on x86 servers [13]. In 3GPP 5G NR related specifications, the above Base Station (BS) components are reorganised into the centralized unit (CU), distributed unit (DU) and radio unit (RU) entities, with their deployment following a flexible topology. CU and DU play the role of BBU, and the RU functions the converting between the signals and radio frequency (RF). A more comprehensive review regarding the details of interfaces and radio evolution is presented in [14]. In the meantime, all the hardware design,

²Can be combined with model-free approaches.

specialized software development and intellectual properties of the RAN-related components are still proprietary.

Network operators expect to obtain decoupled, standardized RAN hardware and open-source operating software to relieve current vendor restrictions. Consequently, the O-RAN alliance was founded in February 2018. Its ambitious mission is to reshape the RAN industry, building future RANs on a foundation of virtualized network elements, white-box hardware, and standardized interfaces. The core principles of O-RAN are intelligence and openness, which will lead the direction beyond 5G and 6G.

Fig. 1 demonstrates one example architecture of O-RAN. O-RAN architecture follows 3GPP architecture and interface specifications, while its NFV is as consistent as possible with European Telecommunications Standards Institute (ETSI). The service management and orchestration (SMO) function in O-RAN has been designed to provide network management functionalities for the RAN and may also be extended to perform core management, transport management, and end-to-end slice management. Meanwhile, the SMO connects with O-Cloud through the O2 interface. O-Cloud is a cloud computing platform comprising a collection of physical infrastructure nodes that meet O-RAN requirements to host the relevant O-RAN functions, the supporting software components and the appropriate management and orchestration functions [2]. One important functionality provided by SMO is the Non-RT RIC designed to implement automated policy-based optimization activities by running ML models. The Non-RT RIC links towards Near-RT RIC via A1. The Near-RT RIC controls and optimizes the functions of CU and DU through the E2 interface. Meanwhile, third-party, microservice architecture-based applications can also be loaded into the Non-RT RIC and Near-RT RIC through rApps and xApps, respectively, to perform data-driven optimization behaviours. In this process, E2 can be leveraged to access the radio node data, and these data can be fed into RICs for ML model training. The CU connects to or controls one or more DUs via the F1 interface. Similarly, one DU connects to at least one RU through the open fronthaul plane. The CU/DU stack hierarchically handles operations of different timescales, while the RU manages and controls the most fundamental RF components and the physical layer in every RU deployment site. All functions of the O-RAN, including the Near-RT RIC, CU, DU and RU, are connected to the SMO through the O1 interface for FCAPS support.

It is noticeable that three control loops involving system parameters and resource allocations are defined in O-RAN. ML solutions can be adopted in any loop based on the time-sensitivity of tasks, in which loop1 handles operators at the time scale of TTI level (<10 ms) for those scenarios that emphasize real-time like the radio resource control and allocation happened in between DU and RU; loop 2 operates in the Near-RT RIC which deals with tasks operating within 10-500 ms. It mainly aims to the O-RAN internal resource control, which RICs perform; loop 3 operates in the Non-RT RIC to process tasks greater than 500 ms.

C. ML/RL APPLICATIONS IN O-RAN

ML is undoubtedly the most remarkable technological progress in recent years. From CV [15], NLP [16] to robotics [17], gaming [18], e-commercial [19] and biology [20] etc. ML applications in almost every technical field have made marvellous achievements. Also, the upcoming O-RAN through the introduction of intelligent programmable RIC enables the RAN to have a mechanism to use emerging learning-based technologies to automate network functions, improve network efficiency, and reduce operating costs. In O-RAN, the initially closed internal radio resources are opened and controlled by unified RICs. That brought some profound changes to communications studies.

- 1) With higher mobile edge computing (MEC) capability, O-RAN enables interaction with end-users, such as directly perceiving end-users behaviours and responding to them so that the optimization of the network can be completed from a more fine-grained and more direct user model analysis way, without the need to perform it in the core network or the centralized cloud.
- 2) O-RAN can significantly help the further promotion of 5G. As often mentioned, the main goals of 5G are enhanced mobile broadband (eMBB), ultra-reliable low-latency communications (URLLC) and massive machine-type communications (mMTC) [21]. Due to the complex and diverse environmental conditions faced by 5G networks, it is necessary to allocate resource blocks with network slicing to meet task requirements for different application scenarios. The introduction of O-RAN makes a dynamic, learning-based slicing mechanism possible. Therefore, deep learning-based adaptive slicing, the collaboration of SDN and NFV, is becoming a research hotspot.
- 3) RICs provide a platform for third-party applications deployment, including ML models, enabling the rapid development and deployment of innovative ideas and algorithms.

The O-RAN use case whitepaper [2] described some of the AI-based deployment targets, such as service level agreement (SLA) assured 5G RAN slice, context-based dynamic handover management for vehicle-to-everything (V2X), traffic steering, and flight path based dynamic unmanned aerial vehicle (UAV) resource allocation etc., while we believe the potential of AI-enabled O-RAN is far more than that. The state-of-the-art communication system embodies a feature of hierarchical and self-contained functions. All functions are interconnected with standardized interfaces. For instance, the signal undergoes a series of units from the transmitter to the receiver, such as modulation, coding, demodulation, denoising, and corresponding channel measurement. Each unit has a well-defined mathematical model that can approach the Shannon limit, and it can be considered that a single unit has achieved its local optimum. However, there are significant challenges in the analysis and optimization of cross-units. If the whole of the above units is regarded as the optimization

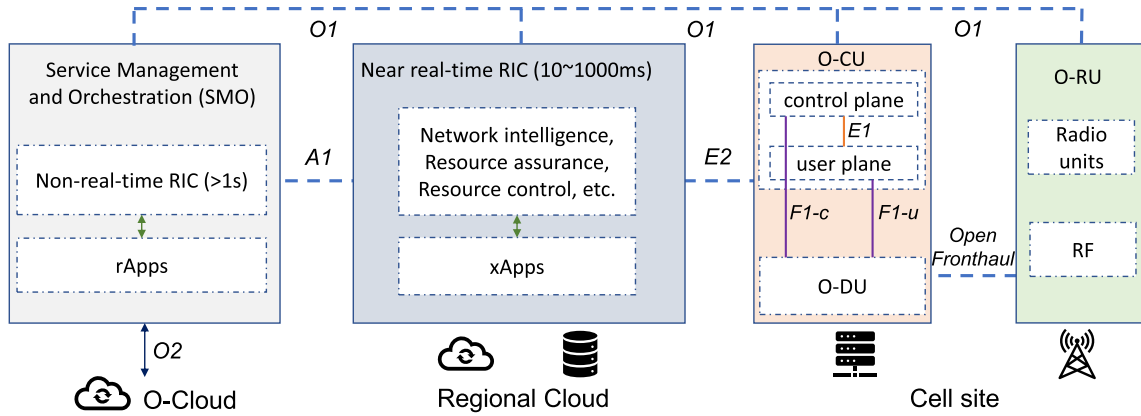


FIGURE 1. One example architecture of O-RAN.

TABLE 2. Survey of the SOTA works regarding DRL applications under the O-RAN context. According to the attributes of these tasks, we divide them into four categories: network slicing, scheduling, and splitting; connection management; resource allocation and xApps related.

Paper	Task	Algorithm	Gains
Abedin et.al. [22]	Elastic O-RAN slicing	Actor-Critic	50% at severed devices
Bonati et.al. [23]	RAN slicing allocation and scheduling	PPO	20% at spectral efficiency
Filali et.al. [24]	RAN resource slicing	Double DQN	Robust and efficient performance for URLLC services
Pamuklu et.al. [25]	Dynamic function splitting	SARSA, Q-learning	Efficiency on renewable energy usage and cost
Polese et.al. [26]	RAN slicing allocation and scheduling	PPO	Improved PRB ration and throughput; smallest buffer occupancy for the MTC traffic; 30% at fewer data requirements of online-training
Lien et.al. [27]	Session management for URLLC	SARSA, Q-learning and double Q-learning	Enabling the gNB to grant a new URLLC session or not
Mollahasani et.al. [28]	Dynamic DU selection	Soft Actor-Critic	50% at energy efficiency
Orhan et.al. [29]	Optimisation of user-cell association	Graph RL	10% - 140% at throughput, cell coverage or load balancing
Wang et.al. [30]	CU–DU resource assignment	Neural MCTS	5.70% - 12.95% at resource utilisation efficiency
Iturria-Rivera et.al. [31]	Power and radio resource allocation	Multi-agent DRL	Higher energy utilization and throughput
Mungari [32]	Radio resource management	-	Dynamic resource allocation based on traffic flow
Zhang et.al. [33]	Power and radio resource allocation	Team DQN	Higher system throughput and lower packet drop rate
Giannopoulos et.al. [34]	Power allocation	Multi-agent DRL	72.3% at energy efficiency
Bonati et.al. [35]	Facilitate the xApps development	-	An integrated platform for data collection, model development and xApps building
Zhang et.al. [36]	Multiple xAPPs coordination	Federated RL	11% higher throughput and 33% lower delay

object, then this kind of global or multi-objective optimization is currently challenging to achieve [37]. The combination with ML/RL various learning paradigms makes O-RAN have the potential for this overall or multi-objective optimization revealed in some advanced research. For instance, in the physical layer, the DL-based OFDM receiver can achieve accurate channel estimation using fewer pilot signals [38]; the end-to-end learning of communication systems has been realized in an autoencoder way which shows advantages in synchronization, equalization and dealing with hardware impairments such as non-linearities [39]; the BS downlink channel state information (CSI) in frequency division duplexing (FDD) massive MIMO system can be inferred by DL with feeding the downlink CSI under certain conditions [40]; under the

premise of imperfect CSI, the design of hybrid massive mimo digital precoder and analog combiner based on RL [41]; and a variety of DL-based LDPC decoding solutions under harsh noise [42]. In the network layer, the learning-based algorithms shape the SON with dynamic resource allocation properties like automated networking, slicing, dynamic spectrum sensing, random access channel, 5G cooperative communication and resource allocation [43], [44], and load balancing optimization in the network layer. It is to be noted that with the O-RAN stepping into the market gradually, some DRL-based optimization cases targeting O-RAN’s features are beginning to appear. A survey of state-of-the-art (SOTA) works regarding DRL applications embracing the O-RAN is shown in Table 2. According to the attributes of these

tasks, we divide them into four categories: network slicing, scheduling, and splitting; connection management; resource allocation and xApps development related. The corresponding algorithms and gains are also detailed in this table.

D. CHALLENGES OF ML/RL DEVELOPING IN O-RAN

Although the hierarchical structure and decoupling characteristics of O-RAN have brought the benefits of supplier diversification, this also brings in higher complexity in O-RAN deployment. On the other hand, developing and deploying a suitable intelligent model in O-RAN may lead to practical engineering technology problems. Considering the most vigorous ML domains in CV and NLP, some standard data sets are generally used to evaluate the performance of the developed ML algorithms. These algorithms are designed to target the features of the given training sample. For example, for the image sample, the initial features of the image space are extracted from adjacent pixels through various convolution operations, and then through various sophisticated network structures such as AlexNet, VGG, and ResNet, the features are further refined. The mapping from the training space to the target space is accurately constructed. In the booming RL field, whether it is gaming or robot control, the development of related algorithms is basically carried out under a standard toolkit like OpenAI gym [65]. A noteworthy phenomenon is that the fields mentioned above benefit from the support of solid mathematical models and complete underlying software. The development of involved ML has been systematically transformed into a near-standard industry. These models corresponding to different application scenarios are well defined, making the goal of algorithm development precise and the whole process controllable.

Turn our attention to the application of ML in O-RAN. The state-of-the-art progress made by the current O-RAN alliance is summarised as follows. (1) The programmable and expandable RIC modules are introduced into the O-RAN architecture, and an interface for data collection within the network is defined. (2) With the clarification of the structure definition, a series of ambitious optimization or control goals for resource, traffic flow, and power consumption have been proposed. (3) The workflow of using SL and RL has been standardized. However, the above progress only reflects the possibility of O-RAN embedded ML in a broad and macro sense. Specific to the realistic implementation of the ML models, we will encounter a rather complicated situation. We further consider issues of ML in O-RAN from algorithms development and deployment angles, respectively. From the view of algorithms development, we summarised the potential issues below:

- 1) In O-RAN, data related to model training is difficult to obtain and process. Even the standard interfaces defined in the O-RAN architecture, such as E2, can access DU, CU and other components to collect information inside the network. This data comes, by default, in raw format and without a schema that is not suitable to be directly consumed by ML/RL algorithms. If we

intend to use this field information to train the model, the cost of data collection will be very high.

- 2) For different optimization goals, the required data for neural network training is heterogeneous. The attributes or patterns of various types of data hiding are elusive. For example, for radio traffic, the data flow as a whole is usually non-Euclidean. In some RU-distributed sites, the data does not meet the characteristics of independent and identically distributed (IID), and some data sets have very strong temporal correlations, while the correlation of other data sets is more reflected in the spatial domain. That will pose challenges to the subsequent data processing methods and feature extraction schemes, affecting the overall neural network structure design.
- 3) Some global optimization problems demonstrate the applicability of RL. That poses other challenges for establishing the connection between O-RAN and RL. These challenges are often not about the RL algorithm itself but how to abstract the problem to be solved into the RL framework and define the RL-related environment, action, state, and reward. For instance, the training issue comes along with high-dimensional state and action spaces; the availability of offline models trained from historical logs; the feasibility of online model training but with limited samples or partial observations; the large reward delay or vanishing in RANs; the complexity of multi-agent RL scheme for optimization problems across multiple RANs [66].
- 4) The RAN is the entrance to the entire wireless network and is closest to UEs. Therefore, the data flow in O-RAN is inevitably directly related to UEs. If we want to use these data streams to train neural network models, new requirements will be put forward for the privacy protection of the UEs and the desensitization of related data.
- 5) O-RAN supports multi-vendor third-party ML/RL applications and increases the complexity of the processes and activities related to the network management plane, which may result in action conflicts in execution, especially when resource allocation is involved. The action coordination ought to be considered in the process of model training [33].

We have introduced that xApps are connected to Near-RT RIC in O-RAN as the host of trained ML models. These trained models are pre-stored in O-cloud and managed by the SMO. However, from the view of model deployment, the above system is not enough to overcome the problems that may arise after the model is deployed in the field. On the one hand, the models obtained by SL were trained by specific data sets. After deploying these mature models, one possible consequence is that the sample data characteristics in the model deployment area are inconsistent with the characteristics of the original training set, which will result in model failure; that is, the expected results cannot be correctly received, as the model can not respond to the input features.

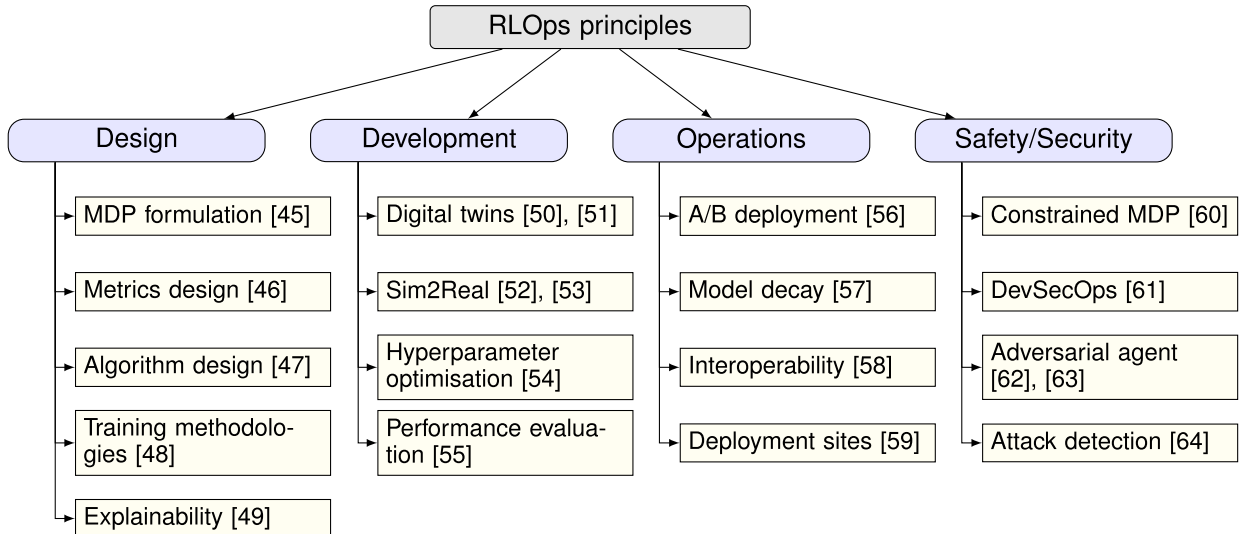


FIGURE 2. This figure lists the critical elements involved in the principle of RLOps. That is the design, development, operations and safety/security, and then we further break down each element into the high-level taxonomy of considerations and methodologies.

On the other hand, as time changes, the external environment changes continuously for the RL model, which will make the initially trained policy no longer suitable. That puts forward new requirements for model management, update and maintenance, and we must look at O-RAN and its ML models from a more holistic perspective.

III. PRINCIPLES OF RLOps

A. BRIEF INTRODUCTION OF MLOps

MLOps is defined as a set of practices that combines ML, DevOps and data engineering, aiming to deploy and maintain ML models in production reliably and efficiently. It can be seen as delivering ML applications through DevOps, with additional attention to data and models. MLOps performs the idea of *automation* and *acceleration*. *Automation* means automating the ML pipeline from data to model for continuous training, as well as automated CI/CD for ML applications. *Acceleration* means to increase the speed of delivery while maintaining the quality of service for ML applications [67].

An MLOps pipeline usually consists of the following elements:

- 1) Data preparation and model design.
- 2) Model testing and validation.
- 3) Model integration, delivery and monitoring.
- 4) Continuous training and CI/CD.

Similar to DevOps, MLOps is an iterative approach. The change in developing requirements, the evolution of the deployment environment, and the alerts raised by monitoring the deployed model would trigger the execution of the pipeline to guarantee the quality of ML applications.

B. MOTIVATION FOR RLOps

MLOps is the general principles and practices of continuous delivery and automation pipelines in ML. Considering the increasing applications of RL in communication networks,

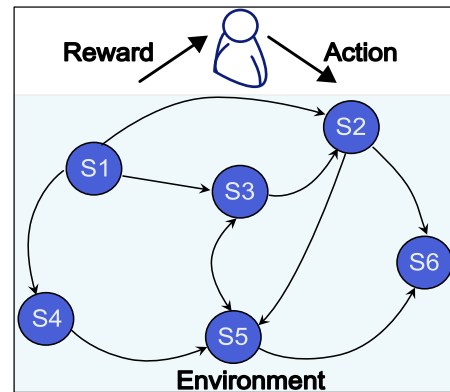


FIGURE 3. Example of RL.

we study the “RLOps” principles to deliver the value of RL to the industry.

RL differs from other ML approaches in several ways, which brings the need for more targeted principal sets. As shown in Fig. 3, *Data & Environment*, *Agent* and *Reward* are the key distinctions considered in the design and delivery of RL applications.

- *Data & Environment*. Data is considered the backbone of ML practices. In RL, data is from agents interacting with environments (online RL) or pre-collected datasets (offline RL) [68]. For online RL, the interaction and learning from live environments (in our case, live communication networks) brings additional risks (as we discussed in Section II-D), which is infeasible in some cases. Hence, the idea of digital twins (DT) has been brought up as a promising solution to the environment and data issue of RL practices [51], providing a controllable, reliable and easily accessible simulation environment. We elaborate it in Section IV-A. Furthermore, considering other real-data hungry cases, communication networks bring challenges to environment access, training data acquisition and model validation.

A network analytics platform with automated data collection, pre-processing, model validation and management abilities are proposed and discussed in Section V.

- *Agent*. Agents are the core of RL problems, interacting with the environment and following their policies. The policy that agents perform is the brain of MDP solutions, as counterparts for “models” in SL and UL. The general principles for developing and deploying ML models also apply to RL models, including model analysis, testing and monitoring.
- *Reward*. The reward is unique to MDPs. It represents the goal of RL, which is essential information to have in model design and deployment. Unlike “labels” as intrinsic features of data in SL, rewards reflect the expected behaviour of agents. In RL applications, reward design is always part of the problem formulation, which requires special attention in RLOps.

As illustrated in Table 2, a large amount of work has been demonstrated for the specific approaches of developing RL-based model in O-RAN recently. However, we believe that on top of these use cases, some common considerations and issues need to be solved, at least, to be realised. By doing that, we expect some general principles of developing and deploying RL models can be summarised to realise true life-cycle management and continuous integration and delivery of such models. Hopefully, a more realistic and affordable RL developing pipeline can be put forward to fulfil the above objective rather than developing case by case. That is the essential intention of this paper.

To effectively deploy RL applications requires careful navigation through a wide range of decisions, from problem formulation to algorithmic choices to the selection of monitoring metrics, to name a few. In an attempt to demystify these decisions, we introduce a non-exhaustive list of “RLOps” principles and observations, which we consider helpful in realizing the potential RL promises. We hope to provide distinct but complementary ideas for RLOps to what may be expected in MLOps and DevOps. For an overview of key considerations and principles for MLOps please refer to [69].

We introduce principles of RLOps under the application development cycle introduced in Fig. 4.³ Below we talk about the three parts: design, development and operation. We will also elaborate on the safety and security concerns related to these three parts. A summary of the high-level taxonomy of considerations and methodologies involved in the RLOps principles is shown in Fig. 2.

C. DESIGN IN RLOps

* *The challenges of design in RLOps lie in the appropriate task formulation and algorithm selection specific to the dynamic environments.*

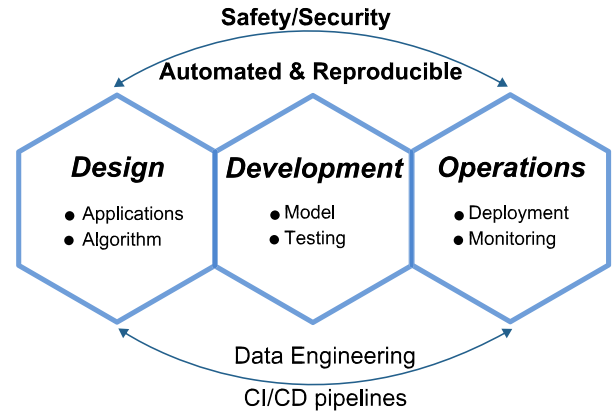


FIGURE 4. RLOps diagram. This diagram demonstrates the development cycle of applications.

1) TASK FORMULATION

Consider the arrival of a new task that takes the form of sequential decision-making, as such RL is likely to be a good solution. Examples of these tasks are given in Section I including handover and interference management, to name but a few [70].

An integral step to build a solution based on RL is to formulate the given problem as an MDP. The formulation of MDP affords many degrees of freedom. Each design decision should be considered carefully, as the form of the MDP will dictate a number of algorithmic decisions. Basic elements to consider include the number of agents, the representation of actions the degree of stochasticity of the environment. For example, if the problem requires distributed decision-making, a stochastic game [71] may be an appropriate formulation; If the hierarchical representation of the action space is possible, and options framework [72] may be possible.

As part of the task formulation phase, it is useful to consider evaluation metrics and baselines that are suitable for the task, where these baselines may be existing solutions. This will smooth out the test, validation and monitoring phases in RLOps, and potentially provide a fail-safe if the RL application begins to behave erratically.

2) ALGORITHM

As discussed in the above section, decisions on the formation of MDP directly impact the form of the solution. Some of the design practices are listed below, but many other general rules exist. Reference [73] provides a good analysis of the impact of some design choices on specific RL algorithms.

- If the action space is continuous, policy gradient-based approaches are likely to be a good option. Some discretisation could also be applicable, such as Q-Learning variants.
- If the state is non-markovian, recurrency can be introduced through stacking previous states [74] or RNN structures like LSTM [75].

³Inspired and adapted from <https://ml-ops.org/img/mlops-loop-en.jpg>

- Particularly, small state-action spaces may be amenable to tabular approaches [76], which provides a higher degree of interpretability over methods that use function approximation.

In the design phase, the training strategies and tricks are also worth considering to tackle problems like *model generalisation* and *training difficulty*. Considering potential varying environments for deploying RL solutions, we are interested in how the model generalizes and as such our training methodologies should reflect this. It has become a consensus in RL research that models trained within limited instantiation of environments do not generalize well [52]. The utilization of methods like “Domain Randomisation” is essential for model generalisation. As for the training difficulty, the utilization of a training curriculum, where a series of increasingly complex tasks are presented to the agent with the intention of easing learning on difficult tasks [77]. Imitation Learning is another approach to ease the learning process, [78] where an agent is pre-trained with the pre-collected dataset containing expert behaviours [68]. In [79], a binary neural network using neuroevolution is presented to simplify the inference model.

In addition to these design choices, we may wish for our algorithm to possess other characteristics. For example, we may want its decisions to be explainable or for it to be aware of its uncertainty with regard to its state. That moves us to other sub-areas of RL researches like *Explainability* [49] and *Bayesian RL* [80], [81] which deals with these concepts that are important from business, legislative or even safety perspectives.

D. DEVELOPMENT IN RLOps

* *The reliable platform and consistent data stream for model training and testing are critical challenges of development in RLOps.*

Once elements of the *design* have reached sufficient maturity levels, steps can be taken towards formally developing the application’s capabilities. This process involves creating the experimental environment (which will likely be based on a DT), model training, and performance optimization.

1) MODEL

The algorithmic approach defined in the Section III-C2 provides the general structure and algorithm for the model. The next step is to develop the necessary code for the agent. To code everything from “scratch” may seem reasonable but may lead to significant engineering expense for limited gain, especially when a wide array of readily available open-source libraries provide high-quality implementations of a range of SOTA algorithms exist.⁴

Training RL applications in a time-efficient and comprehensive manner requires accessibility to a high fidelity

simulator - where a DT will likely be a good fit for this.⁵ If we take a pessimistic viewpoint, the DT or any simulator is an approximation to the real world, and as such, there will be inconsistencies in behaviour that may, in the worst case, lead to testing values being inconsequential as the differences are so profound that the policies are not transferable. This is a *Sim2Real* challenge and is considered in more detail in Section IV-A.

Once an effective algorithmic and simulation approach has been developed to address this challenge, the next major obstacle in the model development process is hyperparameter optimization, which is an arduous and time-consuming process. In the interest of efficient allocation of resources, this process will benefit from automation.

2) TESTING

In the life circle of DevOps, testing is essential to ensure the performance of software systems. Code sanity testing, unit testing and integration testing are commonly used to validate the software iteration. In MLOps [82], the scope of testing extends to data and models. Here we re-consider testing in the context of deep reinforcement learning (DRL) in future O-RAN.

Once a DRL model has been trained, we require functionality within our pipeline to evaluate the model’s capabilities. For trained DRL models, testing should consider multiple model attributes to give a comprehensive evaluation of the models’ performance. Some dimensions are also considered in MLOps, such as the model relevance and accuracy, the robustness to noise, the generalization ability, and ethical considerations [6]. Other challenges are unique to DRL models, for example, the ability of a DRL model to prioritize useful experiences during learning, to choose long-term beneficial actions, to respond to uncertainty, stochasticity, and environmental changes, and avoid unintended behaviour, etc. The testing and validation of DRL models regarding the dimensions mentioned above remain an open question, leaving space for future work. DT might play an important role in the testing procedure since it is an environment in which we have complete control. Manual testing might be required in some use cases. In addition, model interpretability and explainability are of great importance from the perspective of both developer and network service providers, which should be considered during testing. Considering possible network attacks and security challenges, an adversarial attack should also be integrated into the model testing workflow.

E. OPERATIONS IN RLOps

* *The challenges of operations in RLOps lie in the agile and effective monitoring and identification of errors of RL models among different deployment sites.*

Assuming that a model has passed all required testing and validation steps and has been containerized according to

⁴Ray and OpenAI Baselines, to name but a few.

⁵The general structure of which should follow that mandates by OpenAI Gym for consistency with other open-source platforms.

system requirements, the obvious next consideration should be for model deployment and the associated systems required to support and maintain it. This process will include consideration of the deployment location and monitoring with the intention of providing functionalities for continuous improvement.

1) DEPLOYMENT

A fundamental issue (which is discussed at more length in Section III-E2) is that of discrepancies that exist between development and production environments. This problem is likely to be ever-present and difficult to quantify. As such, other safeguards are likely required to mitigate this risk before wide-scale deployment. An example of this could be using software development practices like alpha-beta type deployments to limit the potential impact on end-users whilst getting an empirical measure of application performance.

The environment in which the agents are deployed tends to be highly dynamic, where changes are likely to alter network behaviours. These changes may include internal factors like device configurations and the deployment of other applications or external factors like changes in user behaviour or seasonal phenomena that may affect wireless propagation characteristics. The manifestation of this dynamic environment is a modification to the underlying MDP, and the performance of RL agents will likely degrade accordingly. This issue is one of the concept drifts [57]. The implication of the dynamic nature of the deployment environment is that the performance of deployed models may reduce over time. This general phenomenon is known as *Model Decay* and will be observable through the agent's reception of the reward. This impact can be mediated through periodic re-training if the reward drops below some pre-defined threshold. An alternative approach is to enable online training, but this does come with risks, most notably the requirement for exploration. An additional risk that may arise is non-stationarity [83], which is a consequence of a deployment consisting of multiple RL agents constituting a Multi-Agent system. Non-stationarity arises when multiple agents are learning policies simultaneously, resulting in uncertainty regarding environment behaviours as state transitions are implicitly dependent on other agents.

To enable interoperability on differing base computational platforms all applications will need to be containerised with their associated internal dependencies for deployment with a platform like Kubernetes [58]. A well-defined REST application programming interface (API) will allow for communication of information between entities such that applications can obtain external information that they require for operation and so that monitoring can be performed and decisions can be made pertaining to applications. Communication between disparate systems within the O-RAN architecture naturally raises considerations for model deployment location. By selecting an appropriate location (be that topological or cloud vs edge) and control loops described in Section II-B, application performance and the wider performance of the

network may be improved, where benefits are related to reduced inference time and a reduction in network traffic due to co-location of applications with their dependencies. These decisions may be particularly important for applications that require very low latency for effective operation.

2) MONITORING

Through a collection of Key Performance Indicators (KPIs), the efficacy of an RL agent can be monitored. This information enables decisions pertaining to the application to be made in an informed manner. For example, if an agent is underperforming, it may be desirable to re-train or even replace the agent with an alternative solution.

Monitoring and evaluating RL application performance in the real world is critical to determining whether or not the application is providing benefit, but this is likely to be challenging. Simple measures like cumulative reward can be utilized but are susceptible to issues like reward hacking [84] and do not provide relative measures compared to other methods. The most thorough approach from a network operator's perspective may be to have human oversight of the decisions that agents are making, but this is not scalable and is likely to be problematic as RL agents are often difficult to interpret. Consideration of concepts like *Explainability* [85] is likely to be essential in providing the necessary administrative oversight, which may be necessary from both a risk and governance perspective. The most appropriate strategy is likely to involve an ensemble of methods, including collating a range of metrics that attest to the application's performance characteristics. These measures may include application-specific measures, like throughput and latency for a resource allocation application and include periodic utilization of AB testing to provide a relative measure against well-understood baselines.

In addition to the impact on reward acquisition, changes within the environment in which the RL agents exist may impact the computational performance of the model [69]. Metrics pertaining to model performance, like inference time, throughput, and RAM usage, will be important in identifying transient behaviours.

F. VERSIONING IN RLOps

* *The challenges of versioning in RLOps lie in the synchronization management among the code, model, hyperparameters and developing tools.*

Versioning, or source control, is the practice of tracking and managing changes during development. O-RAN brings the opportunity to use software-based RICs with open interfaces widely. Flexible and fast iteration software development requires careful versioning, and this also applies to RLOps in O-RAN.

1) DATA

The data preparation in RL is different from SL or UL, as it comes from interacting with the environment. For communication network applications, data could come from either

a running network or a DT. Live network data can be stored and versioned by data management tools like DVC,⁶ Pachyderm⁷ or other built-in tools in ML development frameworks. These tools attach version information to datasets. For artificial data generated by a DT, it is more efficient to give snapshots of the DT, including the simulation scenario, the configuration, the random seeds, etc. Given the versioning information of the DT, we should be able to reproduce the same dataset if needed.

2) MODEL

Versioning of the model is vital for controlling the model deployment, especially when facing environment changes or unexpected failures. Since the training pipeline of RL models for O-RAN takes both live network data and DT, it is important to version the training environment and pipeline as well as the model itself to trace back this self-learning approach. This includes the versioning of training configurations, the production environment, and the versioning of DT and network data mentioned in the previous section. The hyperparameters that correspond to each model should also be versioned.

3) CODE

All the production code during development and deployment should be put into versioning. This includes the code to train the RL model, the code for testing and validation, the code for successfully deploying the trained model, and the application code. In addition, as the training of RL in O-RAN uses DT, the code for the DT development and deployment should also be versioned. The DT itself can be seen as a standalone project which requires proper source control [51].

G. SAFETY AND SECURITY IN RLOps

* *The challenges of safety and security in RLOps lie in the robustness assurance of the developed RL models.*

Model safety and operation security are critical for ML/RL applications in O-RAN. The former can be dealt with by introducing safety constraints into the *Design* and *Development* process. We discuss some principles to follow for the operation security, inspired by the DevSecOps [61], which integrates security measures into the DevOps cycle.

For RL models running on wireless networks, *Safety* is important for service assurance as well as avoiding catastrophic performance decay. In the exploratory learning phase, a common approach is to consider potential safety restrictions that exist in the environments, agents, and actions in advance and formalize them into a Constrained MDP (CMDP), which defines a constrained optimization problem as shown in equation 2. A safety policy is expected to achieve by training on the CMDP [66].

$$\max_{\pi \in \Pi} \mathcal{G}(\pi) \quad s.t. \quad C^k(\pi) \leq V_k, k = 1, \dots, K \quad (2)$$

⁶<https://dvc.org/>

⁷<https://www.pachyderm.com/>

where \mathcal{G} is the cumulative discounted reward of a policy π , $C^k(\pi)$ reflects the cumulative cost incurred by constraint k on a given policy π . Specifically, C^k can be defined as $c_k(s, a)$ which represents the possible constraint in terms of state s and action s . Reference [66] presents one solution to the CMDP, which is called Constrained Policy Optimization (CPO). It searches for the policy that maximizes the reward and satisfies the given constraints, i.e., safety requirements. In [86], the sample efficiency in CMDP is further studied in a model-based manner. Robust MDP has also been considered in the scope of CMDP, leading to a robust soft-constrained solution to the Robust-CMDP problems [87].

Security in communication networks protects the integrity of the system, including but not limited to data, applications and user privacy. The open interfaces in O-RAN bring democratised applications but also increase the chance for deployed applications to be attacked. Considering the potential fast and frequent developing circle enabled by the RLOps, security practices should be considered throughout the process. This is the emerging paradigm of DevSecOps, in which some of the security responsibility is downloaded to developers. In RLOps, we make several suggestions in addition to the standard DevSecOps.

Since RL is running in an interactive way to provide intelligent decisions to the system, it is essential to consider the feedback from the environment at the beginning, including the feedback on security. For example, a special state can be designed for the MDP to indicate the sudden change of agent behaviour, which could be a sign of attack. The adversarial agent can be introduced in the RL training to test the robustness against malicious agents [62], [63]. Inspired by [88], *Monitoring* could also play an essential role in integrated security measures. Attack detection techniques like anomaly detection could be applied to enable security practices through monitoring.

IV. BEST PRACTICES OF RLOps

In this section, we discuss some best practices and effective routines for successfully delivering RL applications as the reflections of general principles presented in Section III. We will elaborate on DT's functionalities and critical features, and then discuss the automation and reproducibility engineering in RLOps, respectively.

A. DIGITAL TWINS

A wide range of working definitions of DT exists [50], where we consider the definition by [89] which is that "A digital twin is a digital representation of a physical item or assembly using integrated simulations and service data". The standardization of wireless network DT is still in progress, but it should be able to provide high-fidelity representations of all components of the current live network. This includes the RAN, core network, and characteristics of users and service behaviours among others. Where each component will be modelled through the use of ML models or emulated elements, for example [51]. As discussed in [51], DTs

offer a wide range of benefits for communications networks, including reducing the deployment costs for new services and supporting network automation and optimization.

Within the context of RLOps, DTs are likely to be an integral part of the development pipeline. Enabling training, testing and validation of RL agents in an environment that provides a good approximation to the real world without the associated risks. The key benefits it provides from an RL perspective are enumerated in the list below.

- 1) **Exploration:** RL algorithms require exploration in order to learn about the environment in which they are operating. Exploration, by definition, is risky, as it requires the execution of actions that have potentially unknown outcomes and could, in principle, be unrecoverable [84]. A DT provides a high-fidelity approximation to the real network where a failure is an option, as any damage is inconsequential as it is reversible.
- 2) **Parallelization:** Sample efficiency is a crucial problem within RL, where agents typically take considerable time to train. The utilization of several environments in parallel can reduce the real clock time that an agent takes to converge [90], [91]. Deployment on the real network does not support this functionality.
- 3) **Validation:** When any new component is added (be that physical hardware, an RL agent or some new software function), there is potential for unforeseen negative behaviours to occur. Mitigating these deployment risks is essential from a business perspective. A DT easily accommodates this desired functionality as it allows for simulation and investigation of network response in a wide variety of scenarios. From an RL-specific perspective, it allows for confirmation of the agent’s capacity for reward acquisition and provides functionality to support the interpretability of RL policies more readily.

In addition to the number of compelling arguments for their utilization, certain risks must be realized, especially when the DT modelling can’t reflect the networks’ reality. Within the remainder of this section, we introduce a well-known challenge considered by the RL robotics community, commonly referred to as *Sim-to-Real* [52]. The associated literature is concerned with training within simulation and deployment within the real world and attempts to mitigate risks associated with approximation error between the two systems. Fundamentally, this same desire and challenge will persist within our pipeline and more widely within telecommunications applications. For a comprehensive survey of the area please refer to [52].

B. AUTOMATION

The realization of an automated development process is undoubtedly the critical factor in any type of DevOps. The training procedure needs to be automated in order to save time and labour, expediting the transition from development to production.

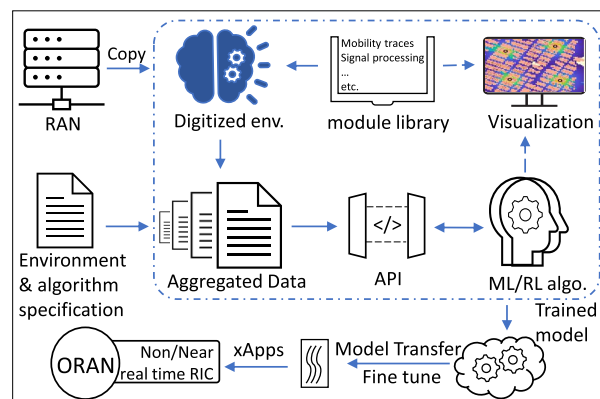


FIGURE 5. Digital twins functionalities. The core, which is a digital copy of the RAN, interacts with additional modules to simulate the change of network environments.

1) DATA

Data cleaning and preparation are often necessary for any new task or environment. This facilitates pattern detection for models as features are well scaled and ordered. As data generated for a task is often consistent, once the transformation procedure is done once, it can be repeated every other time without any need for manual interference. Following data preparation, appropriate feature/state representations must be created to be provided to the agent. This can include concatenating data frames from multiple time-steps together, skipping every n frame, obtaining a certain embedding of the transformed data, etc. This process is often specific to the algorithm/task at hand and is done during training. Since it is a highly repeated step and requires no manual input past creation, it can be automated. The Data transformation and feature extraction process can be finished in the DML and PL layers of our network analytics platform, respectively.

Reward functions can be either extrinsically created for a problem or intrinsically generated from available data. The former case warrants no further automation; however, intrinsic reward signals are obtained from engineering pipelines that extract the signal out of the transformed data. This process will most likely be repeated on every training/evaluation step and must be automated. The data visualisation layer provides such information but needs to establish the automated reward engineering mechanism according to the specific cases.

2) MODEL

Given a certain environment or task, the data preparation pipeline, i.e., the network analytics platform can be triggered and completed automatically. This reduces the amount of time spent on data preparation and guarantees consistency as development evolves. Each RL model follows a specific training methodology. Following data preparation, the training process can also be automated. Training can terminate or resume given performance metrics attached to the agent. At last, for hyperparameter/parameter selection, a common process for DL can also be automated. A hyperparameter sweep can commence once the training pipeline is

formulated. The best set of parameters can be chosen based on performance metrics.

3) CODE

An evaluation/testing step can be automatically triggered once a model has completed training. If passed, the agent can then be deployed to production. This process requires creating rigorous testing scripts, bypassing the agent's manual testing/evaluation, thereby automating the transition from development to production. A model is usually a small sub-part of a larger application infrastructure providing a specific service. Once a new agent is ready for production deployment, it is necessary to automate the new application build process to ensure each new version is well documented and tracked.

C. REPRODUCIBILITY

In O-RAN, well-trained RL models may need to be widely deployed in a large geographic area. Therefore, in the face of different deployment environments and carriers, it is very important to ensure that the performance of the model does not deteriorate, that is the reproducibility.

1) DATA

The development cycle is often about the model, but in many cases can be about changing the environment or handling new data. Changes in data can break a model's performance, and retraining is usually necessary. Dealing with data changes without performance loss is of paramount importance in RL. An agent that can generalize is a flexible and robust one. To tackle the issue of generalization, research challenges have appeared in recent years, such as the Procgen challenge [92] in which agents are tested on multiple versions of the same environment. Keeping track of older versions of data/environments is vital for maintaining stable versions, debugging drops in performance, and developing more robust models.

2) MODEL

Model performance can change drastically with minor changes in the training algorithm. Reproducing results in RL is very difficult given its dynamic nature [93]. In RL both the data source and the agent dynamically change. Moreover, they each influence one another. The environment affects how the agent trains, and the agent's policy impacts the environment's evolution. The ability to revert to stable versions of a model is vital for maintaining stability in the event of performance degradation. In terms of development, minor changes to the model can be researched on their own prior to compounding improvements. Maintaining a careful log of which models possess which mutations are important for ease of integration. Each model version should also contain its own pseudo-code, clearly elaborating the differences in the algorithm. Furthermore, the method of feature creation must be consistent and well logged as it affects how models interpret the provided data. Such strategies massively aid with the development and debugging of new models.

3) CODE

There will be specific dependencies upon which the model relies. Maintaining correct versioning between development and production is necessary for the replication of behaviour. The same goes for the software stack used to create the product in development. It makes no sense to rely on a different, untested stack in production. Therefore, it is often best to containerize development and production iterations. This means all versioning data is well documented within their own containers, allowing for ease of reproducibility.

V. PROPOSED DATA ANALYTICS PLATFORM FOR RLOps

The above sections illustrate the theoretical considerations of RLOps principles and best practices. In order to satisfy the above considerations and to effectively implement ML/RL on top of O-RAN interfaces, a holistic data analytical platform rooting from RAN is necessary, which is helpful for DT continuous refinement, delivers the automation and reproducibility of RL models, and also fulfils the security and confidentiality of multi-tenancy public or private networks. Hence, we design and implement the network analytics platform presented in Fig. 6. We explain the compositions of this platform below.

A. FEATURES OF THE DATA ANALYTICS PLATFORM

- In this platform, the multiple raw data sources need to be collected, validated, enriched, transformed and stored in an integrated data pool. That needs to be processed by data engineering processes, such as application of business rules, creation of KPIs, feature engineering, linkage of data tables according to network topology mapping, etc., which ultimately enables the application of the algorithms according to the targeted use cases.
- Besides, an O-RAN network is built on top of other system components such as IP networks and IT/Cloud infrastructures. The operation and maintenance of these systems are crucial for the overall network performance. It should be integrated into a holistic network management process that addresses all the components.
- Since O-RAN is compliant and allows new architectural models based on multi-tenancy cloudified systems, the data pipeline must guarantee coherence and consistency in the treatment of the different data sources across the whole analytical cycle whilst maintaining strict compliance to the network segmentation and data confidentiality principles guaranteed by the interworking of the data storage, data processing and data governance and policy layers.

B. HIERARCHICAL DEFINITION OF THE DATA ANALYTICS PLATFORM

1) DATA COLLECTION AGENTS

The data collection agents (DCA) are software applications deployed across the network layer, that interact with existing APIs and the network elements (NE). These agents use the standard APIs to collect the standard FCAPS dataset directly

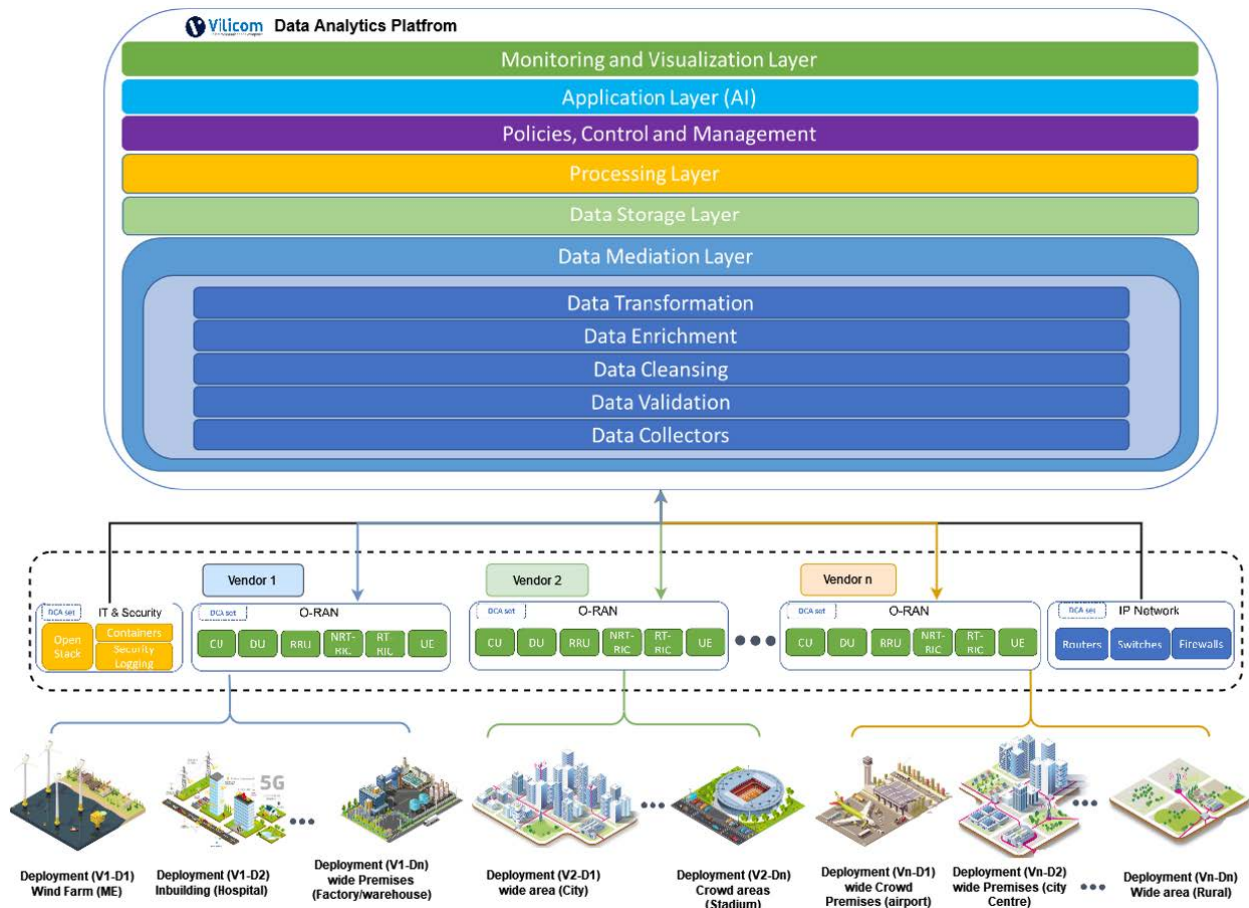


FIGURE 6. This figure illustrates the hierarchical definition and application scenarios of the proposed network analytics platform.

from NEs according to the use case. In a RAN, there are network domains that are implemented using equipment and technology that do not offer open and/or standard APIs. For that reason, it is necessary to develop a specific DCA designed to interact with the specific NE API or protocol, etc. The DCA also has a function of data preparation right from the source, to allow for an efficient and effective data integration coming from multiple and diverse data sources, by normalising the data by applying the conventions that have been defined in the system. The DCA is also responsible for logging all its actions and performing initial data validation procedures. This function is important to trace end-to-end the data pipeline and assist the upper layer of the data mediation stack. These applications are deployed directly on the NE’s management plane or on adjacent servers. These have been designed to listen and track the data generated on these sources and can pull the logs and send them instantly to the data mediation layer (DML).

2) DATA MEDIATION LAYER

The DML is responsible for collecting the data by coordinating the DCAs in the southbound interface, data processing and implementing the northbound interface to the upper layers. This layer is a cluster-based system designed according to big data requirements and best practices [94], allowing

the system to scale and support ultra-dense networks. After data is collected from the DCAs, the DML receives it in its raw format, requiring it to be prepared before going through validation and cleansing processes. The DML needs to add the schema information to the data stream and link it with the network topology. This preparation process increases the efficiency of the system by reducing the complexity of the data validation and data cleansing.

The DML is responsible for the data validation and data cleansing processes that consist in validating the data against the expected schema, identifying duplicate records, or missing records, and coping with latency on the data source in making the data records available. It also prepares the dataset for an optimal application of the data enrichment processes that would fail if applied directly to the raw data due to missing network topology information. The data enrichment and transformation functions are tightly coupled with the data storage and processing layers because it prepares the data stream to match the schemas of the data lake and other consuming applications. At the end of the DML cycle, the data offered to the upper layers are fully integrated, normalised, enriched and transformed according to the system conventions, thus simplifying the development of the data lake and of any processing applications. The DML layers can be continuously improved and extended to consume more

– in quantity and diversity – data sources and to offer the data on the northbound interface in any format, type and frequency that is optimal to the layers consuming the data stream. The DML coordinates with the DCAs to securely collect the data by implementing an encrypted data pipe. It creates one uniform data flow between each DCA and the upper layers.

3) DATA STORAGE LAYER

The data storage layer (DSL) contains one of the main components of the entire architecture which is the data lake. The data lake is the place where the data is stored to be made available to the upper layers, most importantly the processing and application layers. It is designed upon a scalable private cloud object storage; it provides the means to manage and store big datasets that come in diverse formats and structures and enables high throughput and fast access to the data. The policies, business rules, network topology and other meta-data required by the policies, control and management layer are stored in a dedicated relational database that is managed by the DSL. Business Intelligence techniques and the development of ML/RL applications rely heavily upon wide and diverse historical datasets, for trend analysis, statistical analysis and for ML/RL in specific for model training, testing and validation. This demands many computational resources and requires DSL to be designed and implemented using big-data best practices [95], to deliver optimal access to large-scale datasets. On the other hand, feature engineering and RL-related tasks often require high-speed access to many disparate data sources to build and optimise the ML models; this requires high availability of some of the data in great quantities and diversity. For this, we have designed the data lake following the “Cold, Warm and Hot” approach [96].

The data lake is directly accessible by the other layers such as DM, processing and AI layer through a high throughput network. The design behind this storage system allows us to easily store petabytes of data and serve applications regardless of the data access requirements.

4) PROCESSING LAYER

The processing layer (PL) is composed of multiple applications deployed over a containerised environment that scales up with the increased demand from the services of the upper layers, such as the application and visualisation layers. The PL handles mainly three types of jobs, distributed real-time computation, distributed batch processing and jobs related to AI models such as environment states, reward calculation, AI model training/testing, etc. AI and ML applications are complex and hard to develop, maintain, optimise, and deploy because of their iterative and multi-staged life-cycle. Complexity arises mostly from the stages that involve feature engineering, model training, model testing/validation and production deployment. On the other hand, the RL has more components to consider which are the environment, reward calculation, and the agents which make deploying these applications more challenging. As emerged in MLOps

practices, the main enhancement to solve the challenges of the AI lifecycle is to containerize all stages. The PL has been designed and implemented to follow this principle and overcome this challenge. The PL allows the deployment and execution of services that underpin AI applications throughout its life cycle. In addition, to this, it also implements all the services that involve data processing, such as KPI calculation, real-time processing, alarm processing, online monitoring notifications, rule enforcement and data preparation for visualisation. This layer works in tandem with the lower layers, such as DSL and DML, to provide a containerised environment that simplifies the deployment and management of resource-intensive applications and guarantees high-throughput access to the data pool through dedicated and purpose-built data streams. This layer will help to encapsulate the works in subphases where the task could be updated separately without affecting other phases. We illustrate some of the main jobs in this layer as follows:

- **KPI calculation:** To measure the performance of the whole network, 3gpp produced a technical specification document for KPIs [94]. These KPIs need an elevated level of domain expertise to develop and deploy across the data pool. The purpose of these KPIs includes but is not limited to the monitoring and troubleshooting of the network performance and long-term trend analysis of its performance. However, they are valuable features to build ML/RL models and reflect environment status. By abstracting this layer, we intend to save time and reduce complexity. The KPIs are calculated periodically. The results are eventually stored with the AI engineers’ collected performance metrics for usage.
- **Feature engineering and real-time data processing:** Considering the requirements of the RLOps, the processing layer will also run applications that process streams and batches of data, so this layer is where the feature engineering process is done.
- **ML/RL related components:** The components needed to train, test, and validate the ML/RL application. These containers and the related applications are integrated into the whole platform so that they are able to cooperate with other containers and services offered in the processing layer. Additionally, the processing layer is able to run environment simulators or DT images and integrate them into the data pipeline.

5) POLICIES AND CONTROL LAYER

The policies and control Layer is composed of a set of configuration methods, services and metadata that define and implement the business rules, object hierarchy and relationships that are relevant for the functionality implemented across the data mediation, data storage and processing layers. The O-RAN FCAPS data, produced across the multiple virtual network functions (VNFs) and interfaces, is the most representative and important data type in this platform. This data is being structured and is not generated in its raw format, with the whole information that is required for its representation

and integrated with other data sources. This layer contains the rules, metadata, and methodologies necessary for the efficient and effective implementation of the cycles of the DML and DSL, allowing the creation of the structures to validate, cleanse, enrich and store the data in an optimal format. The network topology metadata and methods are fundamental for the linkage of the different managed objects and data structures, thus enabling the cross-layer analysis between network performance events and external events described by data sources - that are external to the O-RAN network and relevant to the analytical process, e.g. UE-based data that describes QoS and QoE events through detailed metrics and logs. On the other hand, this layer also stores the policies and rules that control some aspects of the system's cognitive capabilities, such as the identification of abnormal behaviour and respecting self-healing actions/decisions. These policies and rules can be defined by: subject matter experts (SMEs) through processes of data engineering, feature engineering and/or analytical engineering; and by automated analytic processes, possibly based on ML/RL applications that identify rules/decisions that after being validated and accepted by SMEs are later deployed on to production.

6) AI LAYER (AI APPLICATION MANAGEMENT LAYER)

The AI layer is where the development, initial training and validation of the AI model happens. It allows the implementation of online training through real-time data consumption and offline model validation, generating results/decisions that are not implemented but rather validated by the developers and the subject matter experts. It also allows for monitoring logs and tracking the AI jobs' performance and related application images, mostly for testing and debugging purposes.

7) DATA VISUALIZATION LAYER

This layer is mostly dedicated to implementing business intelligence functions that allow SMEs to access the data in the format of graphical reports and dashboards, thus providing a visual interface to monitor the overall system performance. Through this layer, it is possible to access reports and dashboards that inform about the performance of the different system components through the monitoring of dedicated measurements. The components that are monitored are:

- O-RAN network equipment, VNFs, protocols, interfaces, and functions: this allows for the network management SMEs to evaluate network performance, identify opportunities for optimisation, trends of systemic behaviour and evaluate the impact that AI algorithms might have on the overall system performance.
- AI application decision-making logging: this allows for the DevOps, MLOps and RLOps engineers to evaluate the performance of these applications during the entire life cycle from training to operations. It also allowed to report of the results of correlation and causation analysis visually and emphasised the evaluation of the decision of the application on the system performance.

VI. CONCLUSION


O-RAN embraces the intelligent models in specifications and treats ML/RL as a promising solution for achieving truly intelligent future network infrastructure. Considering the current lack of principles and practices for developing data-derived optimal decision-making strategies in O-RAN, we proposed the RLOps, which takes the life-cycle of RL model development as the main consideration, adopting the design, development, operations and safety/security as principles. We detail all main considerations and methodologies under these principles and integrate the above functions with the digital twins and the network analytics platform, which is geared to achieve automatic and reproducible model operations.

REFERENCES

- [1] L. Bonati, M. Polese, S. D'Oro, S. Basagni, and T. Melodia, "Open, programmable, and virtualized 5G networks: State-of-the-art and the road ahead," *Comput. Netw.*, vol. 182, Dec. 2020, Art. no. 107516.
- [2] C. Li and A. Akman, "O-RAN use cases and deployment scenarios. Towards open and smart RAN," White Paper, 2020.
- [3] M. Soltani, V. Pourahmadi, A. Mirzaei, and H. Sheikhzadeh, "Deep learning-based channel estimation," *IEEE Commun. Lett.*, vol. 23, no. 4, pp. 652–655, Apr. 2019.
- [4] A. Forster, "Machine learning techniques applied to wireless ad-hoc networks: Guide and survey," in *Proc. 3rd Int. Conf. Intell. Sensors, Sensor Netw. Inf.*, 2007, pp. 365–370.
- [5] X. Chen, H. Zhang, C. Wu, S. Mao, Y. Ji, and M. Bennis, "Optimized computation offloading performance in virtual edge computing systems via deep reinforcement learning," *IEEE Internet Things J.*, vol. 6, no. 3, pp. 4005–4018, Jun. 2019.
- [6] C. Renggli, L. Rimanic, N. M. Gurel, B. Karlaa, W. Wu, and C. Zhang, "A data quality-driven view of MLOps," 2021, *arXiv:2102.07750*.
- [7] C. Ebert, G. Gallardo, J. Hernantes, and N. Serrano, "DevOps," *IEEE Softw.*, vol. 33, no. 3, pp. 94–100, 2016.
- [8] V. Dunjko and H. J. Briegel, "Machine learning & artificial intelligence in the quantum domain: A review of recent progress," *Rep. Prog. Phys.*, vol. 81, no. 7, Jul. 2018, Art. no. 074001.
- [9] M. Caron, P. Bojanowski, A. Joulin, and M. Douze, "Deep clustering for unsupervised learning of visual features," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 132–149.
- [10] B. W. Silverman, *Density Estimation for Statistics and Data Analysis*. London, U.K.: Routledge, 2018.
- [11] C. Bishop, *Pattern Recognition and Machine Learning*. New York, NY USA: Springer, 2006.
- [12] K. Zhang, Z. Yang, and T. Basar, "Multi-agent reinforcement learning: A selective overview of theories and algorithms," in *Handbook of Reinforcement Learning and Control*. Cham, Switzerland: Springer 2021.
- [13] Faisal. (Apr. 2021). *Ran Vs Cloud RAN Vs Vran Vs O-RAN: A Simple Guide!*. [Online]. Available: <https://telocloudbridge.com/blog/c-ran-vs-cloud-ran-vs-vran-vs-o-ran/>
- [14] D. Wypiór, M. Klinkowski, and I. Michalski, "Open RAN—Radio access network evolution, benefits and market trends," *Appl. Sci.*, vol. 12, no. 1, p. 408, Jan. 2022.
- [15] N. Akhtar and A. Mian, "Threat of adversarial attacks on deep learning in computer vision: A survey," *IEEE Access*, vol. 6, pp. 14410–14430, 2018.
- [16] D. W. Otter, J. R. Medina, and J. K. Kalita, "A survey of the usages of deep learning for natural language processing," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 2, pp. 604–624, Feb. 2021.
- [17] Y.-H. Wu, Z.-C. Yu, C.-Y. Li, M.-J. He, B. Hua, and Z.-M. Chen, "Reinforcement learning in dual-arm trajectory planning for a free-floating space robot," *Aerosp. Sci. Technol.*, vol. 98, Mar. 2020, Art. no. 105657.
- [18] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, and Y. Chen, "Mastering the game of go without human knowledge," *Nature*, vol. 550, no. 7676, pp. 354–359, 2017.
- [19] H. Yang, "AliGraph: A comprehensive graph neural network platform," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2019, pp. 3165–3166.

- [20] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Židek, A. Potapenko, and A. Bridgland, "Highly accurate protein structure prediction with AlphaFold," *Nature*, vol. 596, no. 7873, pp. 583–589, 2021.
- [21] S.-Y. Lien, K.-C. Chen, and Y. Lin, "Toward ubiquitous massive accesses in 3GPP machine-to-machine communications," *IEEE Commun. Mag.*, vol. 49, no. 4, pp. 66–74, Apr. 2011.
- [22] S. F. Abedin, A. Mahmood, N. H. Tran, Z. Han, and M. Gidlund, "Elastic O-RAN slicing for industrial monitoring and control: A distributed matching game and deep reinforcement learning approach," *IEEE Trans. Veh. Technol.*, vol. 71, no. 10, pp. 10808–10822, Oct. 2022.
- [23] L. Bonati, S. D'Oro, M. Polese, S. Basagni, and T. Melodia, "Intelligence and learning in O-RAN for data-driven NextG cellular networks," *IEEE Commun. Mag.*, vol. 59, no. 10, pp. 21–27, Oct. 2021.
- [24] A. Filali, B. Nour, S. Cherkaoui, and A. Kobbane, "Communication and computation O-RAN resource slicing for URLLC services using deep reinforcement learning," 2022, *arXiv:2202.06439*.
- [25] T. Pamuklu, M. Erol-Kantarci, and C. Ersoy, "Reinforcement learning based dynamic function splitting in disaggregated green open RANs," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2021, pp. 1–6.
- [26] M. Polese, L. Bonati, S. D'Oro, S. Basagni, and T. Melodia, "CoLO-RAN: Developing machine learning-based xApps for open RAN closed-loop control on programmable experimental platforms," *IEEE Trans. Mobile Comput.*, early access, Jul. 4, 2022, doi: [10.1109/TMC.2022.3188013](https://doi.org/10.1109/TMC.2022.3188013).
- [27] S.-Y. Lien, D.-J. Deng, and B.-C. Chang, "Session management for URLLC in 5G open radio access network: A machine learning approach," in *Proc. Int. Wireless Commun. Mobile Comput. (IWCMC)*, Jun. 2021, pp. 2050–2055.
- [28] S. Mollahasani, T. Pamuklu, R. Wilson, and M. Erol-Kantarci, "Energy-aware dynamic DU selection and NF relocation in O-RAN using actor-critic learning," *Sensors*, vol. 22, no. 13, p. 5029, Jul. 2022.
- [29] O. Orhan, V. N. Swamy, T. Tetzlaff, M. Nassar, H. Nikopour, and S. Talwar, "Connection management xAPP for O-RAN RIC: A graph neural network and reinforcement learning approach," in *Proc. 20th IEEE Int. Conf. Mach. Learn. Appl. (ICMLA)*, Dec. 2021, pp. 936–941.
- [30] X. Wang, J. D. Thomas, R. J. Piechocki, S. Kapoor, R. Santos-Rodríguez, and A. Parekh, "Self-play learning strategies for resource assignment in open-RAN networks," *Comput. Netw.*, vol. 206, Apr. 2022, Art. no. 108682.
- [31] P. E. Iturria-Rivera, H. Zhang, H. Zhou, S. Mollahasani, and M. Erol-Kantarci, "Multi-agent team learning in virtualized open radio access networks (O-RAN)," *Sensors*, vol. 22, no. 14, p. 5375, Jul. 2022.
- [32] F. Mungari, "An RL approach for radio resource management in the O-RAN architecture," in *Proc. 18th Annu. IEEE Int. Conf. Sens., Commun., Netw. (SECON)*, Jul. 2021, pp. 1–2.
- [33] H. Zhang, H. Zhou, and M. Erol-Kantarci, "Team learning-based resource allocation for open radio access network (O-RAN)," in *Proc. IEEE Int. Conf. Commun. (ICC)*, May 2022, pp. 4938–4943.
- [34] A. Giannopoulos, S. Spantideas, N. Kapsalis, P. Gkonis, L. Sarakis, C. Capsalis, M. Vecchio, and P. Trakadas, "Supporting intelligence in disaggregated open radio access networks: Architectural principles, AI/ML workflow, and use cases," *IEEE Access*, vol. 10, pp. 39580–39595, 2022.
- [35] L. Bonati, M. Polese, S. D'Oro, S. Basagni, and T. Melodia, "OpenRAN Gym: An open toolbox for data collection and experimentation with AI in O-RAN," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Apr. 2022, pp. 518–523.
- [36] H. Zhang, H. Zhou, and M. Erol-Kantarci, "Federated deep reinforcement learning for resource allocation in O-RAN slicing," 2022, *arXiv:2208.01736*.
- [37] Y. Ouyang, L. Wang, A. Yang, M. Shah, D. Belanger, T. Gao, L. Wei, and Y. Zhang, "The next decade of telecommunications artificial intelligence," 2021, *arXiv:2101.09163*.
- [38] T. Van Luong, Y. Ko, N. A. Vien, D. H. N. Nguyen, and M. Matthaiou, "Deep learning-based detector for OFDM-IM," *IEEE Wireless Commun. Lett.*, vol. 8, no. 4, pp. 1159–1162, Aug. 2019.
- [39] A. Felix, S. Cammerer, S. Dörner, J. Hoydis, and S. T. Brink, "OFDM-autoencoder for end-to-end learning of communications systems," in *Proc. IEEE 19th Int. Workshop Signal Process. Adv. Wireless Commun. (SPAWC)*, Jun. 2018, pp. 1–5.
- [40] Y. Yang, F. Gao, G. Y. Li, and M. Jian, "Deep learning-based downlink channel prediction for FDD massive MIMO system," *IEEE Commun. Lett.*, vol. 23, no. 11, pp. 1994–1998, Nov. 2019.
- [41] Q. Wang, K. Feng, X. Li, and S. Jin, "PrecoderNet: Hybrid beamforming for millimeter wave systems with deep reinforcement learning," *IEEE Wireless Commun. Lett.*, vol. 9, no. 10, pp. 1677–1681, Oct. 2020.
- [42] Y. Wang, Z. Zhang, S. Zhang, S. Cao, and S. Xu, "A unified deep learning based polar-LDPC decoder for 5G communication systems," in *Proc. 10th Int. Conf. Wireless Commun. Signal Process. (WCSP)*, Oct. 2018, pp. 1–6.
- [43] W. Guo, N. M. F. Qureshi, I. F. Siddiqui, and D. R. Shin, "Cooperative communication resource allocation strategies for 5G and beyond networks: A review of architecture, challenges and opportunities," *J. King Saud Univ.-Comput. Inf. Sci.*, Aug. 2022.
- [44] A. K. Bashir, R. Arul, S. Basheer, G. Raja, R. Jayaraman, and N. M. F. Qureshi, "An optimal multitier resource allocation of cloud RAN in 5G using machine learning," *Trans. Emerg. Telecommun. Technol.*, vol. 30, no. 8, p. e3627, Aug. 2019.
- [45] S. Padakandla, "A survey of reinforcement learning algorithms for dynamically varying environments," *ACM Comput. Surveys*, vol. 54, no. 6, pp. 1–25, Jul. 2021.
- [46] S. Jordan, Y. Chandak, D. Cohen, M. Zhang, and P. Thomas, "Evaluating the performance of reinforcement learning algorithms," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 4962–4973.
- [47] K. Arulkumaran, M. P. Deisenroth, M. Brundage, and A. A. Bharath, "Deep reinforcement learning: A brief survey," *IEEE Signal Process. Mag.*, vol. 34, no. 6, pp. 26–38, Nov. 2017.
- [48] T. M. Moerland, J. Broekens, A. Plaat, and C. M. Jonker, "Model-based reinforcement learning: A survey," 2020, *arXiv:2006.16712*.
- [49] E. Puiutta and E. M. Veith, "Explainable reinforcement learning: A survey," in *Proc. Int. Cross-Domain Conf. Mach. Learn. Knowl. Extraction*. Cham, Switzerland: Springer, 2020, pp. 77–95.
- [50] A. Fuller, Z. Fan, C. Day, and C. Barlow, "Digital twin: Enabling technologies, challenges and open research," *IEEE Access*, vol. 8, pp. 108952–108971, 2020.
- [51] H. X. Nguyen, R. Trestian, D. To, and M. Tatipamula, "Digital twin for 5G and beyond," *IEEE Commun. Mag.*, vol. 59, no. 2, pp. 10–15, Feb. 2021.
- [52] W. Zhao, J. P. Queralta, and T. Westerlund, "Sim-to-real transfer in deep reinforcement learning for robotics: A survey," in *Proc. IEEE Symp. Ser. Comput. Intell. (SSCI)*, Dec. 2020, pp. 737–744.
- [53] A. Kadian, J. Truong, A. Gokaslan, A. Clegg, E. Wijmans, S. Lee, M. Savva, S. Chernova, and D. Batra, "Sim2Real predictivity: Does evaluation in simulation predict real-world performance?" *IEEE Robot. Autom. Lett.*, vol. 5, no. 4, pp. 6670–6677, Oct. 2020.
- [54] J. T. Springenberg, A. Klein, S. Falkner, and F. Hutter, "Bayesian optimization with robust Bayesian neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 29, 2016, pp. 4134–4142.
- [55] Y. Li, "Deep reinforcement learning: An overview," 2017, *arXiv:1701.07274*.
- [56] J. Gauci, E. Conti, Y. Liang, K. Virochsiri, Y. He, Z. Kaden, V. Narayanan, X. Ye, Z. Chen, and S. Fujimoto, "Horizon: Facebook's open source applied reinforcement learning platform," 2018, *arXiv:1811.00260*.
- [57] J. Lu, A. Liu, F. Dong, F. Gu, J. Gama, and Z. Zhang, "Learning under concept drift: A review," 2020, *arXiv:2004.05785*.
- [58] V. Medel, O. Rana, J. Á. Bañares, and U. Arronategui, "Modelling performance & resource management in Kubernetes," in *Proc. 9th Int. Conf. Utility Cloud Comput.*, Dec. 2016, pp. 257–262.
- [59] A. Alwarafy, M. Abdallah, B. S. Ciftler, A. Al-Fuqaha, and M. Hamdi, "Deep reinforcement learning for radio resource allocation and management in next generation heterogeneous wireless networks: A survey," 2021, *arXiv:2106.00574*.
- [60] E. Altman, *Constrained Markov Decision Processes*, vol. 7. Boca Raton, FL, USA: CRC Press, 1999.
- [61] H. Myrbacken and R. Colomo-Palacios, "Devsecops: A multivocal literature review," in *Proc. Int. Conf. Softw. Process Improvement Capability Determination*. Cham, Switzerland: Springer, 2017, pp. 17–29.
- [62] T. Chen, J. Liu, Y. Xiang, W. Niu, E. Tong, and Z. Han, "Adversarial attack and defense in reinforcement learning—from AI security view," *Cybersecurity*, vol. 2, no. 1, pp. 1–22, Dec. 2019.
- [63] R. Elderman, L. J. Pater, A. S. Thie, M. M. Drugan, and M. A. Wiering, "Adversarial reinforcement learning in a cyber security simulation," in *Proc. ICAART*, 2017, pp. 559–566.
- [64] A. Pattanaik, Z. Tang, S. Liu, G. Bommanna, and G. Chowdhary, "Robust deep reinforcement learning with adversarial attacks," 2017, *arXiv:1712.03632*.
- [65] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "OpenAI gym," 2016, *arXiv:1606.01540*.

- [66] G. Dulac-Arnold, N. Levine, D. J. Mankowitz, J. Li, C. Paduraru, S. Gowal, and T. Hester, "Challenges of real-world reinforcement learning: Definitions, benchmarks and analysis," *Mach. Learn.*, vol. 110, no. 9, pp. 2419–2468, Sep. 2021.
- [67] D. Sculley, G. Holt, D. Golovin, E. Davydov, T. Phillips, D. Ebner, V. Chaudhary, and M. Young, "Machine learning: The high interest credit card of technical debt," in *Proc. SEML, Softw. Eng. Mach. Learn. (NIPS Workshop)*, 2014, pp. 1–9.
- [68] S. Levine, A. Kumar, G. Tucker, and J. Fu, "Offline reinforcement learning: Tutorial, review, and perspectives on open problems," 2020, *arXiv:2005.01643*.
- [69] E. Breck, S. Cai, E. Nielsen, M. Salib, and D. Sculley, "The ML test score: A rubric for ML production readiness and technical debt reduction," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, Dec. 2017, pp. 1123–1132.
- [70] D. Guo, L. Tang, X. Zhang, and Y.-C. Liang, "Joint optimization of handover control and power allocation based on multi-agent deep reinforcement learning," *IEEE Trans. Veh. Technol.*, vol. 69, no. 11, pp. 13124–13138, Nov. 2020.
- [71] L. S. Shapley, "Stochastic games," *Proc. Nat. Acad. Sci. USA*, vol. 39, no. 10, pp. 1095–1100, Oct. 1953.
- [72] C. Florensa, Y. Duan, and P. Abbeel, "Stochastic neural networks for hierarchical reinforcement learning," in *Proc. Int. Conf. Learn. Represent.*, 2017, pp. 1–17. [Online]. Available: <https://openreview.net/forum?id=B1oK8aoxe>
- [73] N. Rao, E. Aljalbout, A. Sauer, and S. Haddadin, "How to make deep RL work in practice," 2020, *arXiv:2010.13083*.
- [74] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing Atari with deep reinforcement learning," 2013, *arXiv:1312.5602*.
- [75] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [76] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: MIT Press, 2018.
- [77] S. Narvekar and P. Stone, "Learning curriculum policies for reinforcement learning," in *Proc. 18th Int. Conf. Auto. Agents MultiAgent Syst.*, Richland, SC, USA, 2019, pp. 25–33.
- [78] A. Hussein, M. M. Gaber, E. Elyan, and C. Jayne, "Imitation learning: A survey of learning methods," *ACM Comput. Surv.*, vol. 50, no. 2, pp. 1–35, 2017.
- [79] R. Valencia, C.-W. Sham, and O. Simmen, "Using neuroevolved binary neural networks to solve reinforcement learning environments," in *Proc. IEEE Asia Pacific Conf. Circuits Syst. (APCCAS)*, Nov. 2019, pp. 301–304.
- [80] M. Ghavamzadeh, S. Mannor, J. Pineau, and A. Tamar, "Bayesian reinforcement learning: A survey," *Found. Trends Mach. Learn.*, vol. 8, nos. 5–6, pp. 359–483, 2015.
- [81] M. Strens, "A Bayesian framework for reinforcement learning," in *Proc. ICML*, 2000, pp. 943–950.
- [82] *MLOps: Continuous Delivery and Automation Pipelines in Machine Learning*. Accessed: Sep. 20, 2022. [Online]. Available: <https://cloud.google.com/architecture/mlops-continuous-delivery-and-automation-pipelines-in-machine-learning>
- [83] G. Papoudakis, F. Christianos, A. Rahman, and S. V. Albrecht, "Dealing with non-stationarity in multi-agent deep reinforcement learning," 2019, *arXiv:1906.04737*.
- [84] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané, "Concrete problems in AI safety," 2016, *arXiv:1606.06565*.
- [85] E. Puiutta and E. M. Veith, "Explainable reinforcement learning: A Survey," in *Machine Learning and Knowledge Extraction (Lecture Notes in Computer Science)*, A. Holzinger, P. Kieseberg, A. M. Tjoa, and E. Weippl, Eds. Cham, Switzerland: Springer, 2020, pp. 77–95.
- [86] A. HasanzadeZonuzi, A. Bura, D. Kalathil, and S. Shakkottai, "Learning with safety constraints: Sample complexity of reinforcement learning for constrained MDPs," in *Proc. AAAI Conf. Artif. Intell.*, 2021, vol. 35, no. 9, pp. 7667–7674.
- [87] R. H. Russel, M. Benosman, and J. Van Baar, "Robust constrained-MDPs: Soft-constrained robust policy optimization under model uncertainty," 2020, *arXiv:2010.04870*.
- [88] J. Diaz, J. E. Perez, M. A. Lopez-Pena, G. A. Mena, and A. Yague, "Self-service cybersecurity monitoring as enabler for DevSecOps," *IEEE Access*, vol. 7, pp. 100283–100295, 2019.
- [89] R. Vrabic, J. A. Erkoyuncu, P. Butala, and R. Roy, "Digital twins: Understanding the added value of integrated models for through-life engineering services," *Proc. Manuf.*, vol. 16, pp. 139–146, Jan. 2018.
- [90] D. Horgan, J. Quan, D. Budden, G. Barth-Maron, M. Hessel, H. van Hasselt, and D. Silver, "Distributed prioritized experience replay," in *Proc. Int. Conf. Learn. Represent.*, 2018, pp. 1–19. [Online]. Available: <https://openreview.net/forum?id=H1IDy-0Z>
- [91] L. Espeholt, H. Soyer, R. Munos, K. Simonyan, V. Mnih, T. Ward, Y. Doron, V. Firoiu, T. Harley, I. Dunning, S. Legg, and K. Kavukcuoglu, "IMPALA: Scalable distributed deep-RL with importance weighted actor-learner architectures," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 1407–1416.
- [92] S. Mohanty, J. Poonganam, A. Gaidon, A. Kolobov, B. Wulfe, D. Chakraborty, G. Semetulsis, J. Schapke, J. Kubilius, J. Pašukonis, and L. Klimas, "Measuring sample efficiency and generalization in reinforcement learning benchmarks: Neurips 2020 procegen benchmark," in *Proc. NIPS*, vol. 133, H. J. Escalante and K. Hofmann, Eds. Dec. 2021, pp. 361–395. [Online]. Available: <https://proceedings.mlr.press/v133/mohanty21a.html>
- [93] P. Henderson, R. Islam, P. Bachman, J. Pineau, D. Precup, and D. Meger, "Deep reinforcement learning that matters," in *Proc. AAAI Conf. Artif. Intell.*, 2018, vol. 32, no. 1, pp. 1–8.
- [94] Y. Roh, G. Heo, and S. E. Whang, "A survey on data collection for machine learning: A big data—AI integration perspective," *IEEE Trans. Knowl. Data Eng.*, vol. 33, no. 4, pp. 1328–1347, Apr. 2021.
- [95] W. Chang and N. Grady, "NIST big data interoperability framework: Volume 1, definitions," Tech. Rep., Oct. 21, 2019.
- [96] T.-H. Dang-Ha, D. Roverso, and R. Olsson, "Graph of virtual actors (GOVA): A big data analytics architecture for IoT," in *Proc. IEEE Int. Conf. Big Data Smart Comput. (BigComp)*, Feb. 2017, pp. 162–169.



PEIZHENG LI (Graduate Student Member, IEEE) received the B.Eng. degree in optoelectronic information engineering from the Nanjing University of Posts and Telecommunications, China, in 2015, and the M.Sc. degree (Hons.) in image and video communications and signal processing from the University of Bristol, U.K., in 2019, where he is currently pursuing the Ph.D. degree with the Communication Systems and Networks Group. His research interests include machine learning, radio localization, and radio access networks.



JONATHAN THOMAS received the M.Eng. degree in electronics and communications engineering from Cardiff University, in 2016. He is currently pursuing the Ph.D. degree with the Communications Systems and Networks Group, University of Bristol. His research interests include cooperative multi-agent reinforcement learning and its applications in next-generation network management.



XIAOYANG WANG received the B.E. degree in electronic science and technology and the Ph.D. degree in signal and information processing from the University of Electronic Science and Technology of China (UESTC), Chengdu, China, in 2013 and 2018, respectively. She has visited the University of Bristol, Bristol, U.K., from 2017 to 2018. She has joined the Department of Electrical and Electronic Engineering, University of Bristol, as a Research Associate, in 2018.

She has worked on a variety of computer vision topics such as target detection, remote sensing image processing, and visual tracking. Her current research interests include machine learning, especially reinforcement learning, and its applications in next-generation network management.



with self-supervised learning techniques.

AHMED KHALIL received the B.S. degree in mechanical engineering from Purdue University, USA, in 2017, and the M.Sc. degree (Hons.) in robotics from the University of Bristol and the University of the West of England, U.K., in 2019. He is currently pursuing the Ph.D. degree in electrical and electronics engineering with a focus on deep reinforcement learning with the University of Bristol. His current research interest includes bolstering deep reinforcement learning methods



achieved whilst leading data science projects, requiring the design and implementation of end-to-end big data analytical platforms to serve telecom businesses. His research interest includes building (near-RT, RT) platforms that combine O-RAN RAN intelligent controller with reinforcement learning applications.

ABDELRAHIM AHMAD received the dual master's degrees in big data systems and data science and web sciences. He is currently the Data Scientist Lead at Vilicom, where his most important work is related with the development of data and artificial intelligence related applications focused on the telecommunication industry. He has relevant and rich experience in AI, data architecture, big data ecosystems, and 5G and O-RAN technologies. Most remarkable experience have been



experience across multiple areas of expertise such as network management, network operations, network deployment, and network support systems. He has both experience as an engineer and as a technologist that thinks of and about the network life-cycle and also thinks how to develop technology to assist engineers in this life-cycle. His current research interest includes contributing for the development of cognitive networks that are highly automate and autonomic across the whole network management life-cycle.

RUI INACIO received the dual master's degrees in electronics and telecommunications engineering and business and technology management. He is currently a Vilicom's Chief Technology Officer being responsible for the digitization of the company's business models mainly through the development of digital platforms that combine O-RAN technology, big data technologies, and AI technologies. He has built a diversified career in the Mobile Communications Industry, with a rich



machine learning techniques in next generation wireless networks to resource and topology management, connected and autonomous vehicle networks, O-RAN architecture, and cognitive radio networks.

SHIPRA KAPOOR received the Ph.D. degree in electronic engineering from the University of York, U.K., in 2019, with the subject of her thesis being 'Learning and Reasoning Strategies for User Association in Ultra-dense Small Cell Vehicular Networks.' Since 2020, she has been a member of the Self-Learning Networks Team in Applied Research at BT, where she is currently a Research Specialist. Her current research interests include application of artificial intelligence,



automation.

ARJUN PAREKH received the M.Phys. degree in physics from the University of Oxford. He is a Visiting Industrial Fellow at the University of Bristol. He leads a BT Research Team exploring the uses of machine learning and AI for the automation of optimization and planning in converged networks. He led the development of processes and tools to drive the rollout of the U.K.'s first 4G and 5G networks. He has over 20 years experience in radio network management, optimization, and automation.



multiple antenna systems, mmWave communications, and sixth-generation communications systems. She is the author of over 200 journal and conference papers in these areas.

ANGELA DOUFEXI (Member, IEEE) received the B.Sc. degree in physics from the University of Athens, Greece, in 1996, the M.Sc. degree in electronic engineering from Cardiff University, Cardiff, U.K., in 1998, and the Ph.D. degree from the University of Bristol, U.K., in 2002. She is currently a Professor of wireless networks at the University of Bristol. Her research interests include vehicular communications, new waveforms, resource allocation, massive MIMO,



He is currently a R&I Senior Manager at InterDigital, where he is responsible for Radio Interface Research and Development at the Future Wireless Europe Laboratory. He also serves as the Chair of the ETSI Industry Specification Group on Reconfigurable Intelligent Surfaces (ISG RIS) and the Technical Manager of the CELTIC-NEXT European Collaborative Research and Development Project on AI-enabled Massive MIMO (AIMM). Prior to InterDigital, he was a Principal Researcher and the 3GPP RAN1 WG Delegate at British Telecommunications plc. Prior, he held a postdoctoral research positions at University College London, University of Manchester, and King's College London.

ARMAN SHOJAEFARD (Senior Member, IEEE) received the B.Eng. degree in information systems engineering from the Imperial College London, in 2008, and the M.Sc. degree (Hons.) in signal processing and the Ph.D. degree in wireless communications from the King's College London, in 2009 and 2012, respectively. He has extensive experience in project management, research, development, standardization of radio transmission technologies, protocols, and architectures.



connected intelligent technologies and data sciences.

ROBERT J. PIECHOCKI is a Professor of wireless systems at the University of Bristol. His research interests include connected intelligent systems, wireless and self-learning networks, information and communication theory, statistics, and AI. He has published over 200 papers in peer-reviewed international journals and conferences and holds 13 patents in these areas. He leads wireless connectivity and sensing research activities for the IRC SPHERE Project (winner of 2016 World Technology Award). He is a PI/CI for several high-profile projects in networks and AI funded by the industry, EU, Innovate U.K., and EPSRC, such as NG-CDI, AIMM, OPERA, FLOURISH, and SYNERGIA. He regularly advises the industry and the Government on many aspects related to connected intelligent technologies and data sciences.

...