



Empirical performance analysis and ML-based modeling of 5G non-standalone networks

Konstantinos Kousias^{a,*}, Mohammad Rajiullah^b, Giuseppe Caso^b, Ozgu Alay^{a,b}, Anna Brunstrom^b, Usman Ali^c, Luca De Nardis^c, Marco Neri^d, Maria-Gabriella Di Benedetto^c

^a University of Oslo, Norway

^b Karlstad University, Sweden

^c Sapienza University of Rome, Italy

^d Rohde & Schwarz, Italy

ARTICLE INFO

Keywords:

Mobile networks

5G NSA

Machine learning

ABSTRACT

Fifth Generation (5G) networks are becoming the norm in the global telecommunications industry, and Mobile Network Operators (MNOs) are currently deploying 5G alongside their existing Fourth Generation (4G) networks. In this paper, we present results and insights from our large-scale measurement study on commercial 5G Non Standalone (NSA) deployments in a European country. We leverage the collected dataset, which covers two MNOs in Rome, Italy, to study network deployment and radio coverage aspects, and explore the performance of two use cases related to enhanced Mobile Broadband (eMBB) and Ultra-Reliable Low Latency Communication (URLLC). We further leverage a machine learning (ML)-based approach to model the Dual Connectivity (DC) feature enabled by 5G NSA. Our data-driven analysis shows that 5G NSA can provide higher downlink throughput and slightly lower latency compared to 4G. However, performance is influenced by several factors, including propagation conditions, system configurations, and handovers, ultimately highlighting the need for further system optimization. Moreover, by casting the DC modeling problem into a classification problem, we compare four supervised ML algorithms and show that a high model accuracy (up to 99%) can be achieved, in particular, when several radio coverage indicators from both access networks are used as input. Finally, we conduct analyses towards aiding the explainability of the ML models.

1. Introduction and background

5G cellular systems are developed to address high bandwidth, low latency, and massive connectivity requirements of enhanced Mobile Broadband (eMBB), Ultra-Reliable Low Latency Communication (URLLC), and massive Machine-Type Communication (mMTC) use cases. Therefore, compared to 4G technologies, i.e., LTE and LTE-A,¹ 5G systems are expected to satisfy more demanding and heterogeneous Quality of Service (QoS) and Quality of Experience (QoE) requirements, e.g., in terms of throughput, latency, and reliability.

To support such requirements, the 3rd Generation Partnership Project (3GPP) has standardized two main 5G deployment modes in Release 15 (Rel-15), i.e., NSA and SA [1]. The main difference is that, in its so-called LTE-assisted options, i.e., Options 3, 3a, and 3x, 5G NSA relies on the existing 4G core network for control plane management, while 5G

SA uses its own 5G core, hence, acting independently of the 4G network (Fig. 1). Both modes require a 5G New Radio (NR) RAN composed of gNBs, i.e., the 5G equivalent of 4G eNBs, and corresponding Physical Cell IDs (PCIs), i.e., a reusable cell identifier at the physical layer, which can operate in low (<1 GHz), mid (1–6 GHz), or high (>24 GHz) frequency bands.

In the ongoing 5G deployment stages, the majority of MNOs decided to adopt the NSA mode, since it comprises a more straightforward and less costly solution. However, 4G-5G co-dependence in NSA introduces several challenges in terms of configurations, procedures, and performance, that require further study and investigation. Therefore, it is key to understand the implications of the NSA architecture on coverage, performance, and network management (e.g., gNB placement strategies, handovers, resource allocation, network configurations, and

* Corresponding author.

E-mail addresses: konstako@uio.no (K. Kousias), mohammad.rajiullah@kau.se (M. Rajiullah), giuseppe.caso@kau.se (G. Caso), ozgua@ifi.uio.no (O. Alay), anna.brunstrom@kau.se (A. Brunstrom), usman.ali@uniroma1.it (U. Ali), luca.denardis@uniroma1.it (L. De Nardis), marco.neri@rohde-schwarz.com (M. Neri), mariagabriella.dibenedetto@uniroma1.it (M.-G. Di Benedetto).

¹ For the sake of brevity, Fourth Generation (4G) Long-Term Evolution (LTE) and LTE Advanced (LTE-A) will be referred to as 4G for the rest of this paper.

<https://doi.org/10.1016/j.comnet.2024.110207>

Received 31 August 2023; Received in revised form 12 January 2024; Accepted 21 January 2024

Available online 23 January 2024

1389-1286/© 2024 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

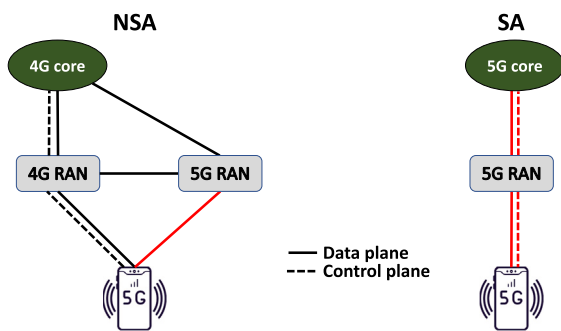


Fig. 1. A schematic diagram of 5G NSA and 5G SA architectures.

so on). In this context, the rapid deployment of 5G NSA on top of 4G networks allows for conducting empirical investigations on the above aspects, towards revealing how deployment, configuration, and technology features affect end user performance (Section 2 provides an overview of such studies). Such data-driven analyses are increasingly important for 5G systems, as they perfectly complement the popular approach of using artificial intelligence (AI) and ML for 5G and beyond-5G network management and optimization, for which the collection and use of reliable datasets is required towards proper testing and validation [2].

Among several new features introduced in Rel-15, two main technologies are used in 5G for improving performance, and thus require specific empirical analyses: SSB beamforming, which is available for both 5G NSA and SA architectures, and DC between LTE and 5G, which is specific to 5G NSA Option 3, thus also referred to as EUTRA-NR DC (EN-DC).

On the one hand, SSB beamforming is a technique that uses multiple propagation beams (up to 8 in mid-band, as per Rel-15) to direct the 5G signal towards different UEs; SSB beam (re-)selection depends on the encountered network coverage conditions, and is critical for increasing spectral efficiency and coverage, particularly at high frequency, e.g., millimeter wave (mmWave) [3]. On the other hand, DC allows a UE to be simultaneously connected to and exchange data with 4G and 5G PCIs, aiming to increase data rate and reliability [4,5]. After presenting our measurement setup and campaign in Section 3, our analyses of the impact of SSB beamforming and DC, among other aspects, on 5G NSA coverage and performance are reported in Sections 4 and 5, respectively.

Focusing on DC, in order to exploit the advantages of using this technique, a key aspect to address in current 5G NSA deployments is the definition of proper policies under which a 5G PCI is added as a data carrier for a UE. MNOs can take into account several aspects to define such policies, e.g., the readiness of their 5G deployments, the congestion level on both access networks, and the actual need for a 5G carrier with respect to the traffic demand of the UE. Another key aspect in such decision is the radio coverage experienced by the UE when both 4G and 5G PCIs are available, which the UE reports to the network for aiding the DC decision. This aspect is extensively analyzed in Section 6, where we adopt a ML-based approach to model the DC decision process followed by MNOs in 5G NSA. In the analysis, we leverage the coverage characteristics of the nearby 4G and 5G cells as recorded in multiple locations during our measurement campaign. Such analysis can be utilized towards providing a better understanding of the criteria (in terms of coverage) used to decide the DC mode. In addition, the output models can be used by MNOs and service providers for improving deployment and network configuration aspects in view of the future 5G NSA/SA implementations. Another aspect where the modeling of the DC mode can be exploited is for building DC connectivity maps. A number of different platforms, such as the *OpenCellID* [6], *RTR Netztest* [7] and *LTE Italy* [8] use crowdsourced measurements to

collect the GPS location of a wide number of cell towers from multiple technologies (e.g., 4G, 5G) towards building large scale coverage maps, which in general provide information on the signal quality. However, such approaches do not offer further information about the actual type of connectivity a UE is likely to experience at any given location on the map (e.g., 4G or 5G). We believe that our ML models can be used in conjunction with the existing coverage maps solutions to provide a more rounded view of the cellular ecosystem.

In this paper, we extend our previous contributions [9] and present a large-scale measurement study on the commercial 5G mid-band NSA networks of two major MNOs in Rome, Italy. In particular, we empirically study deployment, coverage, and performance aspects of 5G NSA deployments in a European country, and provide key insights on SSB beamforming and DC technologies. In terms of QoS/QoE, we study throughput and latency/reliability aspects, aiming to showcase the performance achievable by 5G NSA for eMBB/URLLC-like services in comparison with 4G. Moreover, we use supervised ML to model the DC decision process in 5G NSA, aiming to understand whether the addition or removal of a 5G carrier can be accurately predicted by considering the radio coverage characteristics of nearby 4G and 5G cells. The work takes advantage of the dataset introduced in [9] and thoroughly described in [2], where we also provide more information on the dataset, its availability under an open-source license, and its potential use cases.

We make the following contributions:

- Considering three different mobility scenarios, we present a comparative analysis of radio coverage and network deployment strategies between two MNOs in Italy. In addition, we study aspects related to SSB beamforming and complement its impact on coverage.
- We explore the performance of 4G and 5G NSA by focusing on throughput and latency/reliability aspects, executing an extensive analysis across different scenarios.
- We model the DC decision process in 5G NSA using a ML-based methodology. By performing an extensive study of ML algorithms and input features, we show that a high modeling accuracy can be achieved by considering the radio coverage characteristics of both access networks, while we perform additional analyses to address model explainability aspects. Such analysis can be used by MNOs towards improving different network aspects but also for building and integrating DC connectivity maps on top of existing coverage maps.

We organize the paper as follows. Section 2 presents the related work, while Section 3 describes the experimental design and provides details on the collected dataset. Section 4 demonstrates results on network deployment and radio coverage, while Section 5 analyzes QoS/QoE performance. Section 6 focuses on the modeling of the DC decision process using supervised ML. Last, Section 7 concludes the paper.

2. Related work

In recent years, numerous studies based on measurements of 5G systems have surfaced. Apart for initial performance evaluations conducted on dedicated testbeds, e.g., those within the scope of the EU 5G Public Private Partnership (5G-PPP) [10] and the US Platforms for Advanced Wireless Research (PAWR) [11] initiatives, the analysis of 5G performance has been predominantly carried out on commercial 5G networks deployed in urban contexts.

Performance baseline metrics for commercial 5G networks were outlined in distinct studies, i.e., [12] in the US, [13] in China, and [9] in Europe (Italy), where several aspects were investigated, including coverage, throughput, and latency for mid-band and/or high-band deployments, and NSA and/or SA modes, depending on the MNOs' deployments in the regions under study.

Table 1

Dataset statistics per scenario and grouped by MNO, application (active dataset), and RAT/session type. We report the number of samples and sessions for the passive and active datasets, respectively. Note that the passive dataset contains additional measurements (beyond the active counterpart measurements) that we use for a more rigorous statistical analysis.

RAT	Passive (samples)				Active (sessions)											
	4G		5G		4G				5G partial		5G					
	App		App		Speedtest		Gaming		Speedtest		Gaming		Speedtest		Gaming	
MNO	Op_1	Op_2	Op_1	Op_2	Op_1	Op_2	Op_1	Op_2	Op_1	Op_2	Op_1	Op_2	Op_1	Op_2	Op_1	Op_2
IS	89.6K	127K	553K	3.8M	75	49	68	107	50	21	6	9	4	45	79	48
OW	112K	114.8K	456K	2.51M	38	147	264	176	4	18	13	14	1	6	37	33
OD	41.1K	42.9K	156.5K	1.02M	10	31	93	39	10	27	21	7	1	14	76	32

In particular, the authors in [12] conducted a measurement study on 5G mid-band and mmWave networks in the US, thus establishing a performance baseline for early commercial 5G deployments. In [14], they extended their work by introducing a ML-based framework for 5G throughput prediction leveraging gradient boosting and *seq-2-seq* models. In addition, in [15] they analyzed handover management and power consumption aspects, while in [16] they complemented their study by proposing a holistic solution for 4G/5G handover prediction. In [17] they addressed the performance of volumetric video streaming applications in the context of 5G throughput, while in [18], they proposed a crowdsourced platform for data collection. Likewise, the authors in [13] analyzed commercial 5G mid-band networks in China, focusing on coverage, throughput, latency, and device energy consumption aspects. Finally, our previous work in [9] focused on the analysis of coverage, deployment, and performance (throughput and latency) aspects for 5G NSA mid-band networks in Italy, and was complemented by further work in [19–21], where handover implications on performance [19] and outdoor-to-indoor propagation in the mid-band frequency [20,21] were discussed.

In this paper, we further extend [9] with novel analyses. In terms of coverage, deployment, and performance, we provide additional insights on the impact of different network configurations on radio coverage and end-user performance (e.g., uplink throughput). Furthermore, we include a new analysis on the modeling of the DC mode in 5G NSA networks using supervised ML classification. This study showcases an additional ML use case that departs from previous ML use cases in the context of data-driven 5G system analysis, which are mostly related to 5G throughput prediction [14,22], channel propagation modeling [23], and handover prediction/forecasting [16]. We highlight that our analysis differs from the ones related to handovers, as it focuses on understanding how the radio coverage characteristics of both 4G and 5G access networks can be used for the efficient modeling of the DC mode at any given measurement location towards providing further insights on deployment and network configuration aspects.

3. Experimental design and dataset

This section presents the measurement campaign and provides a description of the collected dataset.

3.1. Measurement setup and configuration

We performed a large-scale measurement campaign for two major Italian MNOs, denoted as Op_i ($i \in [1, 2]$) in the following, offering both 4G and 5G coverage in Rome, Italy, during a period of 7 weeks between December 2020 and January 2021.

Measurement Scenarios. The campaign was organized in *sub-campaigns* carried out during different day types (weekdays vs. weekends) and times of the day. The sub-campaigns were conducted in different urban areas in Rome, encompassing Municipalities I, VII, VIII, IX, XI, XII, and XIII, with a population density ranging from 1000 (Municipality IX) to 8000 (Municipality I) inhabitants per squared km. Sub-campaigns were carried out under different mobility scenarios, defined as: *IS*, for measurements performed at the 7th floor of a residential




Fig. 2. The measurement system consisting of (a) the R&S TSMA6 toolkit (scanner and PC), (b) a GPS antenna, (c) a commercial 5G-enabled device, (d) an RF antenna, and (e) a tablet for monitoring the PC (in the TSMA6 toolkit) where the ROMES software was running.

building, and different locations (offices and laboratories) at the 2nd floor of the Department of Information Engineering, Electronics and Telecommunications (DIET) of Sapienza University of Rome; *OW*, for measurements performed outdoor while walking with an average speed of 3.63 km/h; and *Outdoor Driving (OD)*, for measurements performed outdoor while driving a private vehicle (car) with an average speed of 17.36 km/h. In both mobile scenarios, the average speed is significantly affected by the encountered traffic and road conditions (e.g., traffic lights, road congestion, etc.).

Measurement System. The campaign was carried out using the R&S TSMA6 toolkit, an integrated system composed by a network/spectrum scanner and an embedded Intel PC running Windows. The PC was used to run the R&S ROMES software, which allowed for real-time visual inspection, campaign replay, and data storing/exporting. As shown in Fig. 2, the TSMA6 toolkit was equipped with (i) a synchronized GPS antenna to enable the geo-mapping of the locations where measurements were collected, (ii) a commercial 5G-capable device (Samsung S20) for performing active measurements, and (iii) an RF omnidirectional antenna to allow for the continuous passive scanning of downlink control signals from operational 3GPP technologies. The scanner was configured to perform passive measurements on four 4G bands, i.e., Bands 1 (2110-2170 MHz), 3 (1805-1880 MHz), 7 (2620-2690 MHz), and 20 (791-821 MHz), and one 5G band, i.e., Band n78 (3.3–3.8 GHz). We used the R&S Android application called *Qualipoc* for running the active measurements, i.e., QoS/QoE tests. Note that our setup reports measurements on the millisecond granularity. In addition, for the active measurements, an event-based measurement update is used, thus the reported data may be not evenly spaced in time.

3.2. Measurement methodology

To evaluate the performance of the networks under test, we used both passive and active measurements.

Passive Measurements. We used the scanner to detect and decode downlink control information, aiming to infer network deployment and coverage characteristics. Due to the high scanner sensitivity, the collected data is representative of the coverage experienced by the UE in the same area, but also includes information on signals that a normal UE may not be able to detect and decode (e.g., signals from distant PCIs), thus allowing for a finer analysis.

To represent each measurement, we use the notation $m_{lat, long}^{freq}$, where $[lat, long]$ and $freq$ are the estimated geo-location coordinates pair and the operating carrier frequency, respectively. For 4G, $k \geq 1$ measurement samples were recorded at each location, where k represents the number of detected 4G PCIs in that location. For 5G, due to the possible use of SSB beamforming, measurements were reported at SSB beam level. Hence, for each 5G PCI, the scanner decoded the characteristics from $1 \leq l \leq 8$ SSB beams.

Active Measurements. We performed active measurements through a 5G-capable UE (Fig. 2), configuring it to perform QoS/QoE throughput and delay/reliability performance tests. Such tests were repeated multiple times during each sub-campaign. We use the term *session* to refer to each repetition.

We adopted two UE modes: (i) *5G-enabled*: the UE exposes its 5G NSA capability so that the network can decide whether to connect it to 5G PCIs along with 4G PCIs, thus making DC possible, and (ii) *5G-disabled*: the UE only exposes its 4G capability so that the network only connects it to 4G PCIs (DC is thus not possible in this case). Therefore, sessions were partitioned as (i) *5G*, when a 5G connection was available throughout the session, (ii) *5G Partial*, when 5G was partially available during the session (with automatic fallback on 4G), and (iii) *4G*, when only 4G was available. These latter sessions either belong to 5G-disabled campaigns or to 5G-enabled campaigns where a 5G PCI was never added as a further data carrier for the UE (e.g., due to 5G coverage holes). Next, we provide a brief description for each active test.

- **Throughput Test.** We used the Speedtest tool by Ookla [24] to measure the end-to-end throughput in both 5G-enabled and 5G-disabled modes, with a target server located in Rome. The length of each session is of about 60 seconds. Speedtest uses Transmission Control Protocol (TCP) for performing throughput measurements, in single or multiple connections modes. We used the second option where a proprietary algorithm is used for determining the number of connections [25].
- **Latency/Reliability Test.** We performed the Interactivity test provided by R&S Qualipoc [26] to measure latency and reliability. The methodology adopted in this test has been taken by European Telecommunications Standards Institute (ETSI) as the base for defining 5G performance testing procedures [27], and has also been approved by the International Telecommunication Union Telecommunication Standardization Sector (ITU-T) as a standard methodology [28].

In the test, the UE acts as a client and sends packets to a server, which in turn reflects each packet back to the UE. The adopted transport protocol is User Datagram Protocol (UDP), while the higher layer protocol is Two-way Active Measurement Protocol (TWAMP) [29]. The rationale of the test is that a service-dependent interactivity score (*i-score* [%]) can be defined as the perceived responsiveness of a service, and evaluated as a function of three main Key Performance Indicators (KPIs): (i) latency, defined as the RTT of a sent-and-reflected packet, (ii) PDV, defined as the difference between the latency of a packet and the minimum latency across all the packets, as per RFC 5481 [30], and (iii) packet loss, defined in terms of missed packet percentage, i.e., the ratio between packets that are lost in the network or are received after a target delay budget and the total number of packets scheduled at the UE side. The *i-score* is then evaluated by using a model that takes into account the specificity of each service. More details on *i-score* modeling can be found in [27,28].




Fig. 3. Traffic pattern used for emulating real-time online gaming services.

For our tests, we used the real-time online gaming traffic pattern shown in Fig. 3. This pattern is symmetric in DL and UL, and emulates different phases of a typical online multi-player gaming application, with low-to-medium data rate (from 0.1 to 1 Mbps). The test duration is 10 seconds and the two-way delay budget is 100 ms, according to 3GPP [31]. We executed all the tests towards a server located in Switzerland.

3.3. Dataset statistics

Table 1 shows statistical information of our dataset for each scenario, MNO, application (for active measurements), and Radio Access Technology (RAT) vs. session type for passive vs. active measurements.

Passive dataset. We used a subset of the available features, including spatial and temporal fields, carrier frequency identifiers, and signal strength/quality indicators, i.e., RSRP, RSRQ, and SINR. For 4G, such indicators were measured on the Reference Signal (RS) sent by 4G PCIs; for 5G, they were measured on the Secondary Synchronization Signal (SSS), sent by 5G PCIs in the SSBs [20,21]. Hence, for 5G, the dataset includes SS-RSRP, SS-RSRQ, and SS-SINR at SSB beam level.²

The 4G and 5G passive datasets consist of approximately 527K and 8.48M samples, respectively (see Table 1). The 5G dataset size differs significantly between the two MNOs since, during our measurements, Op_1 was using a single SSB, while Op_2 adopted SSB beamforming with up to 8 SSBs.

Active dataset. We conducted 555 Ookla Speedtest sessions (i.e., 197 for Op_1 and 358 for Op_2) and 1158 real-time online gaming sessions (i.e., 657 for Op_1 and 501 for Op_2). Table 1 shows the number of sessions for different scenarios, MNOs, and session types.

4. Network deployment and radio coverage

5G NSA network deployment presents several challenges due to 4G/5G co-dependence and adopted carrier frequencies, among others, which may lead to reduced radio coverage and possible connectivity issues. Therefore, it is important to understand the deployment choices of the MNOs and study radio coverage characteristics in different scenarios.

4.1. Network deployment

In this sub-section, we present network deployment statistics for each MNO and provide insights on the adopted SSB beamforming strategies.

² The open-sourced dataset also provides additional features that cover a variety of 5G signals, i.e., Demodulation (DM), Physical Broadcast Channel (PBCH), Primary Synchronization Signal (PSS), and SS-PBCH.




Fig. 4. Spatial representation of the RAN deployment for Op_1 . The blue and red markers indicate the estimated positions of 4G eNBs and 5G gNBs, respectively. The black line in the background indicates the UE routes during the entire measurement campaign. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4.1.1. Deployment density

Across all sub-campaigns, the scanner reported a total number of 313 4G and 162 5G PCIs for Op_1 , while, the corresponding values for Op_2 are 407 and 158, respectively. In addition, the number of detected eNBs for Op_1 and Op_2 are 125 and 117, respectively. The above numbers show that Op_2 has a denser 4G PCI deployment compared to Op_1 , while the 5G deployment is rather similar for the two MNOs. At the time of our campaign, MNOs were collectively using around 30% of the available PCIs for 5G, hinting at an under way 5G deployment.

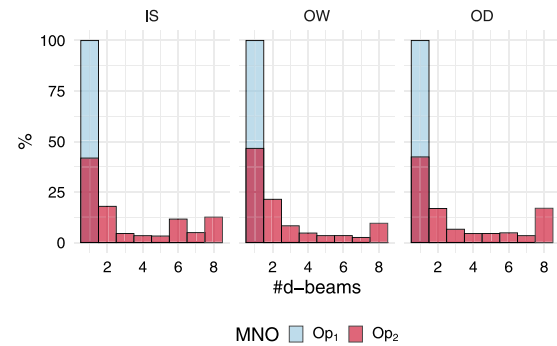
As a reference example, Fig. 4 shows the eNB/gNB spatial distribution for Op_1 (similar results were obtained for Op_2). We use different colors to differentiate between the two RATs, i.e., blue for 4G eNBs and red for 5G gNBs. The black line represents the UE OW and OD routes during the entire measurement campaign. Note that PCI position estimation is a feature provided by ROMES. To infer the position of each eNB/gNB, we perform additional processing (merging) of the data. In addition, for 5G, the position of the gNBs is approximated using the position estimates of the 5G PCIs and the nearest 4G eNB. This is a reasonable assumption considering that the NSA control plane is executed by the 4G core. Therefore, the 5G RAN is in most cases co-located with the 4G RAN, as also confirmed by online and freely-accessible databases, e.g., LTE Italy [8].

Overall, the results reveal that, at the time of the measurement campaign, the 5G RAN was still sparse compared to its 4G counterpart, possibly resulting in 5G coverage holes.

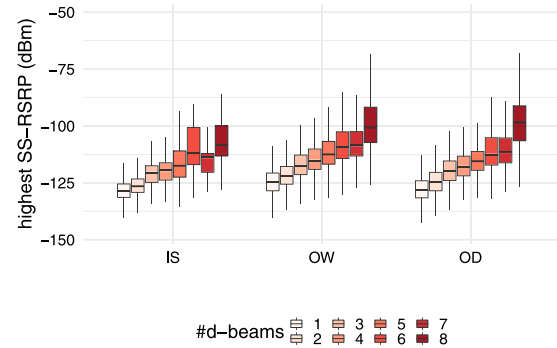
4.1.2. SSB beamforming strategies

Fig. 5(a) depicts the density histogram of $\#d\text{-beams}$, defined as the number of detected SSBs per PCI at each location, grouped by scenarios and colored by MNO.

We observe that, at the time of the measurement campaign, Op_1 did not support SSB beamforming, hence, the percentage for $\#d\text{-beams} = 1$ across all scenarios is 100%. On the contrary, the scanner detected up to 8 SSB beams for the PCIs of Op_2 . In this case, we observe that the likelihood of detecting a single SSB beam is significantly higher than the rest $\#d\text{-beams}$ counts. This is due to the possible presence of locations with insufficient 5G coverage, where the high scanner sensitivity allows for detecting spurious SSBs from likely distant PCIs. Furthermore, for all scenarios, we see that the likelihood of detecting higher SSB beam counts becomes small with the exception of $\#d\text{-beams} = 8$. We believe this is due to the presence of locations with excellent 5G coverage (e.g., the scanner was in Line of Sight (LoS) with a given PCI), where although SSBs are beamformed, the scanner was still able to detect



(a) Density plot of $\#d\text{-beams}$ colored by MNO.



(b) Boxplot distribution of the highest SS-RSRP SSB beam across different $\#d\text{-beams}$ counts (Op_2).

Fig. 5. SSB beamforming deployment analysis. Both subfigures are grouped by scenario. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

and decode all the beams transmitted by surrounding PCIs (e.g., due to signal reflection). Across scenarios, we did not notice significant differences. Therefore, a 5G UE would not be able to detect and decode all the available SSBs from PCIs in urban scenarios unless it operates under good coverage conditions.

Fig. 5(b) complements the above analysis by showing the impact of SSB beamforming on coverage, by depicting the highest perceived SS-RSRP at each $\#d\text{-beams}$ count for Op_2 . For each scenario, the median values follow an increasing trend as $\#d\text{-beams}$ increases, i.e., the probability of a higher SS-RSRP increases when multiple SSB beams are detected. We further look into the average number of 5Gs PCIs for Op_2 uniquely detected at each measurement point during our sub-campaigns. This number varies among the different scenarios, i.e., 11 for IS, 11 for OW, and 7 for OD. Even though less PCIs are reported in the OD scenario, we observe that, on average, SS-RSRP in OW and OD is higher than SS-RSRP in IS by approximately 3 dB. This result indicates that a higher density of deployed and detected PCIs does not always result in a higher coverage, as it is possible that such PCIs are either far away or are not in LoS with the UE (e.g., due to buildings or other obstacles). A single, stronger (in terms of RSRP) PCI would be sufficient to achieve a better coverage, particularly when outdoor and indoor scenarios are compared, considering the additional power losses due to signal penetration in the indoor case [20]. We highlight that these initial discussions on coverage are in line with the results shown in the next sub-section, where we provide further details on both MNOs.

Takeaway remarks: The above results provide useful insights for better understanding the current deployment strategies followed by the operators while pinpointing significant shortcomings, such as the 5G coverage holes (shown in Fig. 4). In addition, they reveal to what extent 5G SSB beamforming is currently adopted by each operator as

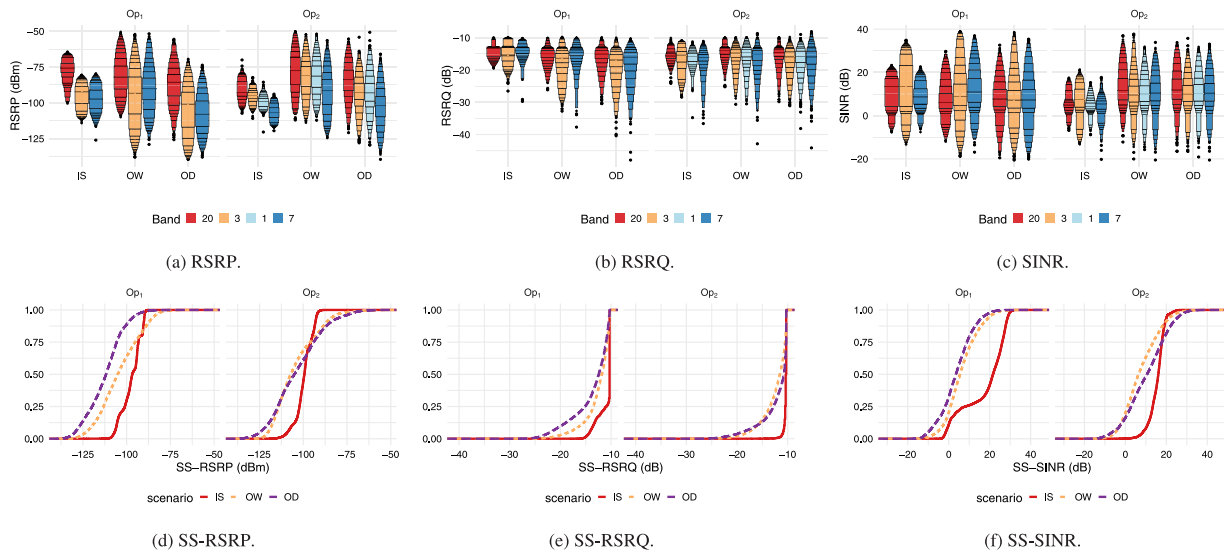


Fig. 6. Distribution of (a) 4G RSRP [dBm], (b) 4G RSRQ [dB], (c) 4G SINR [dB] across carrier frequencies (in a letter-value plot format and order by highest to lowest frequency) and (d) 5G SS-RSRP [dBm] (e) 5G SS-RSRQ [dB] (f) 5G SS-SINR [dB] (in an ecdf format), for Op_1 and Op_2 and grouped by scenario. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 2

Median values for 4G RSRP [dBm] and 5G SS-RSRP [dBm] for each combination of operator, scenario, and band of operation.

		IS				OW				OD			
Band		B20	B3	B1	B7	B20	B3	B1	B7	B20	B3	B1	B7
4G	Op_1	-75.7	-92.28	-	-97.3	-81.99	-93.27	-	-90.07	-85.57	-107.95	-	-100.95
	Op_2	-92.79	-94.97	-99.12	-106.8	-77.46	-81.01	-81.29	-91.51	-83.73	-92.69	-92.98	-101.97
5G	Band	n78				n78				n78			
	Op_1	-96.8				-105.1				-111.7			
	Op_2	-99.3				-107.2				-105.1			

well as its impact on coverage. MNOs can leverage such results towards deriving strategies for optimal 5G RAN planning and deployment.

4.2. Radio coverage

We now focus on the analysis of radio coverage, showing how it varies between MNOs and across different scenarios.

Given a location, radio coverage is determined as the highest RSRP (4G) and SS-RSRP (5G) perceived among all detected PCIs. To establish a performance baseline, we first look at 4G radio coverage. Figs. 6(a), 6(b), and 6(c) depict the distribution of the observed 4G RSRP, RSRQ, and SINR, respectively, in a letter-value plot format grouped by scenarios, dissected by MNOs, and colored by 4G carrier frequencies. In the following, the analysis will focus only on the RSRP indicator, since all three 4G parameters has a similar behavior across the three scenarios. The upper part of Table 2 summarizes the results for RSRP by presenting the median values for each combination of MNO, scenario and band of operation.

We observe that the two MNOs have quite different coverage properties. In particular, for Op_1 , the median RSRP across all bands for IS and OW scenarios is on par (-90.51 dBm and -89.05 dBm, respectively), while it decreases to -98.84 dBm for OD. On the contrary, Op_2 experiences the best vs. worst RSRP conditions in OW (median value of -82.89 dBm) vs. IS (median value of -97.5 dBm), respectively, while the median RSRP is -92.92 dBm in OD. We also notice that, for both MNOs, the RSRP standard deviation increases between IS and OW/OD, with an average difference of 4 dB and 4.5 dB for Op_1 and Op_2 , respectively, which is expected considering that a mobile UE will experience higher coverage variability. For both MNOs the median RSRP shows a slow degradation as the UE moving speed increases, with an average RSRP difference between OW and OD of 9.79 dB and 10.03 dB for Op_1

and Op_2 , respectively. Last, across the available carrier frequencies, we observe a decreasing pattern for RSRP as the frequency increases, due to worse penetration and propagation properties of high frequency signals.

Likewise, Figs. 6(d), 6(e), 6(f) depict the distribution of the 5G SS-RSRP, SS-RSRQ, and SS-SINR as an empirical CDF function. As both MNOs use a single frequency band for providing 5G service (mid-band), data is no further dissected. To provide a fair comparison between 4G and 5G, we focus our attention on the SS-RSRP indicator; results are summarized in the lower part of Table 2 by presenting the median values for each combination of operator and scenario. We observe that for both MNOs SS-RSRP is significantly higher in IS than in OW/OD scenarios, where drops of up to 11.6 and 7.9 dB are observed for Op_1 and Op_2 , respectively. Interestingly, the results for the two MNOs are rather different when comparing OW and OD scenarios: for Op_1 the coverage is markedly better in OW than in OD, with a gap of 6.6 dB, while Op_2 measurements show similar coverage for the two scenarios, with a slight advantage (2.1 dB) in OD; a likely explanation for the differences between the two MNOs ties back to the use of SSB beamforming by Op_2 , and further highlights its effect on coverage.

Takeaway remarks: On par with the discussions in the previous subsection, these results provide valuable insights on the performance of 5G mid-band radio coverage across different scenarios. Such discussions can be further exploited towards efficient deployment and management of the 5G NSA infrastructure, as well as a baseline for future 5G SA deployments.

5. Application performance

In this section, we present a comparative analysis of throughput and latency/reliability performance between 4G and 5G NSA networks.

Table 3

Median values for DL and UL throughput, and i-score measured during our measurement campaign, for IS, OW, and OD scenarios.

Scenario	Session connectivity	MNO	DL Thput [Mbps]	UL Thput [Mbps]	i-score [%]
IS	4G	Op_1	36.30	8.740	65.50
		Op_2	31.95	22.800	61.40
	5G partial	Op_1	500.00	13.55	65.10
		Op_2	382.00	19.60	85.30
	5G	Op_1	618.50	13.90	70.70
		Op_2	455.00	23.15	86.50
OW	4G	Op_1	27.15	23.95	58.40
		Op_2	20.70	31.60	65.20
	5G partial	Op_1	368.00	13.20	52.20
		Op_2	127.50	14.35	44.40
	5G	Op_1	527.00	18.80	71.40
		Op_2	254.50	10.70	46.00
OD	4G	Op_1	37.40	26.00	59.55
		Op_2	27.30	24.10	67.60
	5G partial	Op_1	249.50	25.55	58.35
		Op_2	190.00	17.15	59.90
	5G	Op_1	372.00	11.30	75.75
		Op_2	340.50	17.10	83.40




Fig. 7. DL throughput grouped by MNO and scenario. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

5.1. Throughput analysis

We start our analysis by presenting the DL and UL throughput results.

5.1.1. Downlink throughput

Fig. 7 shows the achieved DL throughput from Speedtest sessions in different scenarios, grouped by session type and MNO, with a summary of the achieved median DL throughput also provided in Table 3.

We observe that, in comparison to 4G, 5G (i.e., 5G and 5G Partial) provides significant DL throughput gains. In IS, the 5G average median throughput values for Op_1 and Op_2 are 560 and 419 Mbps, respectively, with Op_1 providing the highest DL throughput value of 701 Mbps in a particular 5G Partial session. In contrast, the 4G median throughput values for Op_1 and Op_2 are only 36 and 32 Mbps, respectively.

The 5G throughput gain in IS reduces as the propagation conditions change in OW/OD, with the highest throughput variability experienced in 5G Partial sessions. These sessions include, by definition, situations with unstable DC, i.e., a 5G carrier was added or removed on top of 4G carriers during the execution of the active test session, which we also refer in the following to as handovers from 4G to 5G, and vice versa. Therefore, these results suggest that unstable DC in NSA, e.g., due to poor 5G coverage, can be detrimental to throughput performance, as also observed in [12,13]. It is interesting to observe that for both 5G and 5G Partial sessions the trends and the relative gaps for throughput across scenario follow closely those observed for SS-RSRP, and reported in Table 2. For both MNOs the highest throughput is in fact measured in the IS scenario, where both experience the highest SS-RSRP, but while

for Op_1 the throughput decreases steadily as the mobility increases (as SS-RSRP does), this is not true for Op_2 , for which the throughput in the OD scenario is slightly larger than in the OW one; this is in line with the slightly higher SS-RSRP observed in OD vs. OW.

In order to further analyze throughput performance, Fig. 8 shows examples of time series collected during a 5G-enabled (a) and a 5G-disabled (b) IS campaign for Op_1 . The first sub-figure of Fig. 8(a) shows that the 5G DL throughput achieves about 700 Mbps in most of the sessions. We also observe that 5G uses a MCS index up to 27 (second sub-figure) and consequently a modulation order up to 256-QAM, which is possible in radio conditions as good as the ones registered in these sessions (see the high SINR in the fourth sub-figure). Hence, as reported in the third sub-figure, 5G ends up using a high Transport Block Size (TBS) of around 50 000 Bytes, as a result of the adoption of high MCS and possibly large allocated bandwidth, i.e., the resource blocks allocated to the UE. Assuming the UE is connected to a single PCI/carrier, the total available channel bandwidth is between 5–20 MHz for 4G systems, while it can reach up to 100 MHz for 5G mid-band networks [32]. In our measurements, we observed that both MNOs allocate a channel bandwidth of 78.120 MHz for 5G (resulting from subtracting 940 kHz of each guard band from a 80 MHz spectrum allocation, as per Italian regulation) as opposed to a maximum of 45 MHz for 4G as a result of varying level of Carrier Aggregation (CA) on 4G. When such a large TBS is transported in a short Transmission Time Interval (TTI) of 0.5 ms (due to the sub-carrier spacing of 30 kHz adopted by both MNOs, as confirmed by our measurements), 5G achieves high throughput.

In the 5G-disabled scenario (Fig. 8(b)), 4G shows lower throughput compared to 5G (around 100 Mbps, see first sub-figure), due to the smaller allocated channel bandwidth in 4G, the use of lower MCS index (<20) (second sub-figure) and thus modulation orders. Hence, 4G uses a lower TBS (third sub-figure), which in turn explains the lower throughput. Good radio conditions were also observed for 4G (fourth sub-figure), but to a lower extent compared to 5G, and thus impacting the choice of MCS and modulation.

As mentioned in Section 1, 5G NSA enables DC between 4G and 5G RATs [4]. In our dataset, we observe that, via DC, the UE is connected to both 4G and 5G PCIs. As a result, Op_1 schedules DL traffic on both 4G and 5G carriers, leading to significant capacity contributions on both 4G PDSCHs and 5G PDSCHs. Fig. 9(a) shows a set of time series recorded during an IS campaign for Op_1 , where we see that the DL throughput at the application layer (first sub-figure) includes the contributions from both 5G and 4G PDSCHs (second and third sub-figures). Fig. 9(b) shows the results in the same location for Op_2 , where we see that, when DC is activated, Op_2 predominantly uses the 5G PDSCHs, therefore, resulting in lower throughput.




Fig. 8. An example of DL throughput for 5G (a) and 4G (b), with corresponding values of MCS, TBS, and SINR (Op_1 , IS scenario).




Fig. 9. An example of different 5G and 4G PDSCH aggregation strategies for Op_1 (a) and Op_2 (b) in a IS scenario.

5.1.2. Uplink throughput

We now discuss the achieved UL throughput in different scenarios and for different session types and MNOs.

Fig. 10 shows the UL throughput for each scenario and MNO, with a summary of the achieved median UL throughput also provided in Table 3. First, we observe in the figure that the 5G UL throughput is significantly lower than the DL counterpart (Fig. 7). This is justified by several factors, including UE power limitations and the use of the Time Division Duplex (TDD) scheme in the 5G mid-band under analysis. As indicated in our dataset, the adopted TDD pattern results in a DL/UL resource split of 7:2 slots, which significantly limits the UL opportunities. We also notice reduced MCS and modulation (up to 64-QAM), resulting in smaller TBS and decreased throughput. An example is shown in Fig. 11 for Op_1 in an IS scenario.

When comparing across scenarios, we see that, in IS, the 5G median UL throughput is approximately 5 Mbps higher than 4G for Op_1 , while these two values are quite similar for Op_2 , with a higher variability registered for 4G. Surprisingly, for Op_1 , we observe that both 4G and 5G throughput in IS is lower than the throughput measured in the mobile outdoor scenarios (OW/OD). We explain this result by further inspecting both the usage of DC in UL and the 4G/5G coverage experienced by the UE during the corresponding measurement campaigns. First, we

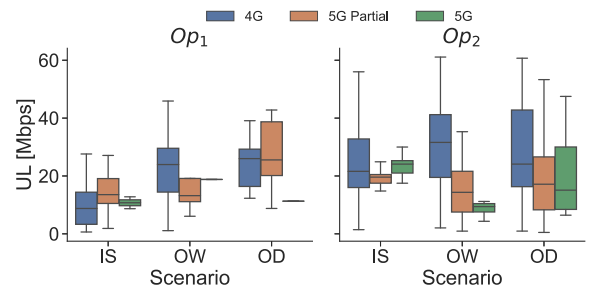


Fig. 10. UL throughput grouped by MNO and scenario. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

notice that, when the UE is connected to both 4G and 5G PCIs via DC, a large portion of the UL traffic is redirected on the 4G carrier, with only a small portion of it actually being transmitted on the 5G carrier via 5G PUSCH (Fig. 12(a) for Op_1). This contrasts with the use of DC in DL for Op_1 (Fig. 9(a)), where we observe that the 5G PCI carries most of the DL traffic. For Op_2 , we instead see a predominant use of the 5G




Fig. 11. An example of achieved UL throughput for 5G with corresponding values of MCS, TBS, and SINR (Op_1 , IS scenario).

PUSCH for UL data transmission (Fig. 12(b)), in alignment with the DL case (Fig. 9(b)).

Second, we observe that the UE experiences the worst average 4G coverage in IS compared to OW/OD, particularly in terms on RSRP (a similar 5G coverage is instead observed across the scenarios). These two aspects together lead to the conclusion that, for Op_1 , the 4G coverage has a significant impact on both 4G and 5G UL throughput due to the pronounced use of the 4G carrier in both cases, with lower 4G coverage leading to lower throughput, as observed for Op_1 in IS. In this same scenario (Op_1 , IS), it is worth observing how, although not fully exploited, adding the 5G carrier in DC mode helps to slightly improve the performance, with 5G *Partial* and 5G sessions having higher throughput than 4G sessions. In addition, in Fig. 10, similarly to the DL case, 5G performance seems more susceptible to mobility (OW/OD), likely due to coverage holes and variability. This results in throughput values lower than 4G, particularly for Op_2 . Overall, our results are in line with previous analyses [12] and suggest that MNOs have focused on assuring significant DL throughput gains for the initial roll out of their 5G NSA networks, with the UL counterpart requiring further enhancement.

Takeaway remarks: Our analysis compares the performance of 4G and 5G NSA networks, highlighting 5G's superior downlink throughput which is a key requirement for eMBB use cases. However, 5G's uplink throughput is notably lower due to factors like UE power limitations and the TDD scheme used. Handovers between 4G and 5G due to, for instance, variable 5G coverage can negatively impact performance. While significant gains in downlink throughput have been achieved with the rollout of 5G NSA networks, there is a need for further improvements in uplink capabilities and handling mobility challenges.

5.2. Latency/reliability analysis

Beyond delivering high throughput, 5G is also expected to lower the communication latency towards supporting URLLC. In this section, we present the latency/reliability performance results obtained during the real-time online gaming tests.

Figs. 13 and 14 show the obtained results in different scenarios, for 4G, 5G *Partial*, and 5G sessions, with a summary of the results with respect to the *i-score* also provided in Table 3.

Figs. 13(a) and 13(b) show that 5G provides the lowest median RTT and PDV across all scenarios and MNOs. However, due to unstable 5G

coverage, a high variability is observed for these KPIs in 5G *Partial* sessions, e.g., in OW/OD for Op_1 . Fig. 14(c) shows the missed packet percentage. Starting with Op_1 , there are almost no missed packets in either 4G or 5G sessions in all scenarios. However, there are missed packets in both mobile scenarios during 5G *Partial* sessions. This clearly shows that, similarly to the throughput case, 5G NSA handovers from 4G to 5G, and vice versa, may lead to latency/reliability performance degradation. Fig. 14(d) (top) shows the *i-score* for Op_1 . Having the lowest median RTT and PDV, and nearly no missed packets, 5G sessions provide the highest *i-score* in all scenarios. However, while 5G *Partial* sessions provide an *i-score* similar to 4G in IS, they also provide the worst *i-score* in OW/OD due to increased missed packets. Considering Op_2 and referencing the bottom of Fig. 14(c), we observe almost no packet loss in IS across all sessions. Yet, in OW/OD, both 5G *Partial* and 5G sessions experience missed packets. This is further illustrated by the bottom of Fig. 14(d), where 5G sessions have the highest *i-score* in IS, while a clear performance decrease for 5G and 5G *Partial* sessions is observed in OW/OD.

As also discussed in [19], these results are tightly coupled with DC policies in 5G NSA. Indeed, if not carefully configured, such policies can lead to sub-optimal choices in adding/removing/using a 5G data carrier, eventually leading to performance losses that appear particularly detrimental for latency/reliability-sensitive services.

Takeaway remarks: Our study highlights that while 5G outperforms 4G in stationary settings for real-time online gaming (a key URLLC use case), its performance in mobile scenarios is less consistent. This inconsistency is attributed to factors like handover policies in NSA and unstable coverage. The conclusion emphasizes that while 5G shows promise in enhancing the latency/reliability for sensitive applications like online gaming, its full potential is not yet fully realized under conditions involving mobility.

6. DC decision process modeling

In this section, we study the modeling of the DC mode in 5G NSA networks leveraging a ML-based methodology. First, we describe the post-processing of our dataset and carry out an exploratory data analysis to gain insights about the dataset's attributes towards understanding whether modeling of the DC mode can be achieved by only leveraging traditional statistics. Next, we present the ML experimental setup, carry out the performance evaluation, and conduct additional analyses focusing on model explainability aspects.

6.1. Dataset post-processing

For modeling the DC mode, we merge both active and passive features (see Section 3) into a unified dataset. In particular, we extract two groups of variables, i.e., (i) RSRP, RSRQ, SINR, SS-RSRP, SS-RSRQ, and SS-SINR and (ii) DC Mode (*5G-Connected* or *5G-NotConnected*, defined in Section 6.3). As the scanner reports the radio coverage properties from multiple 4G/5G cells at each location, we only include the ones offering the highest RSRP/SS-RSRP values, respectively. We further collect additional metadata such as time, UE's geo-location coordinates (Latitude and Longitude), MNO name, and sub-campaign name. Such information is essential for assigning primary keys during merging but also for specific analyses presented in later subsections.

We only consider sub-campaigns that involve a certain degree of mobility, i.e., OW and OD. In particular, we use data from 71 sub-campaigns (40 OW and 31 OD) resulting in approximately 244K measurement samples (i.e., 96K for Op_1 and 148K for Op_2). Since data are not equally distributed between the two DC modes (33%–67% in favor of *5G-NotConnected*), we adopt the Synthetic Minority Oversampling TEchnique (SMOTE) [33] to generate synthetic samples from the minority class (*5G-Connected*). As a result, the sizes of the newly balanced datasets for Op_1 and Op_2 are 125K and 195K, respectively.




Fig. 12. An example of different 5G and 4G PUSCH aggregation strategies for Op_1 (a) and Op_2 (b) in a IS scenario.

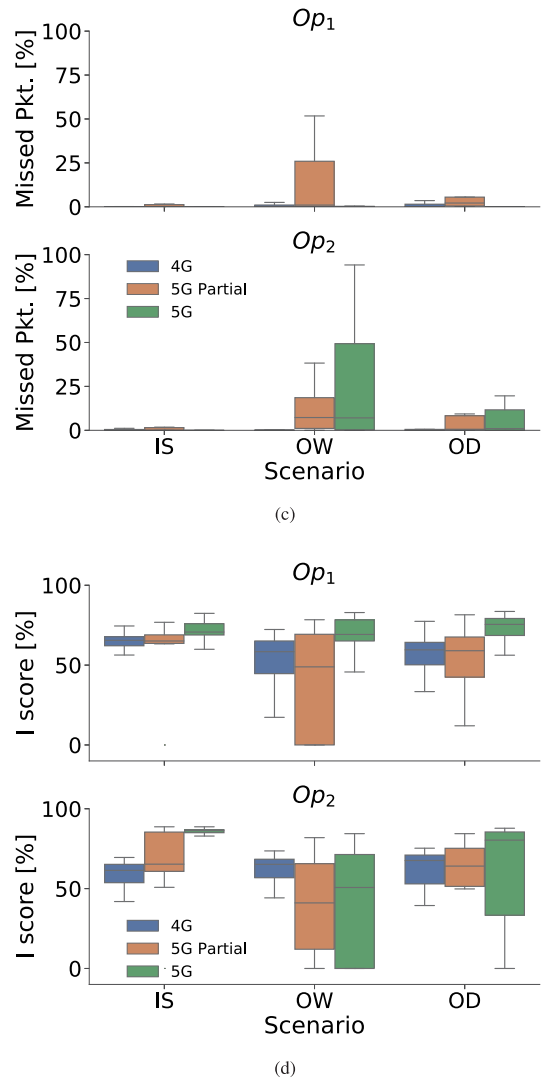
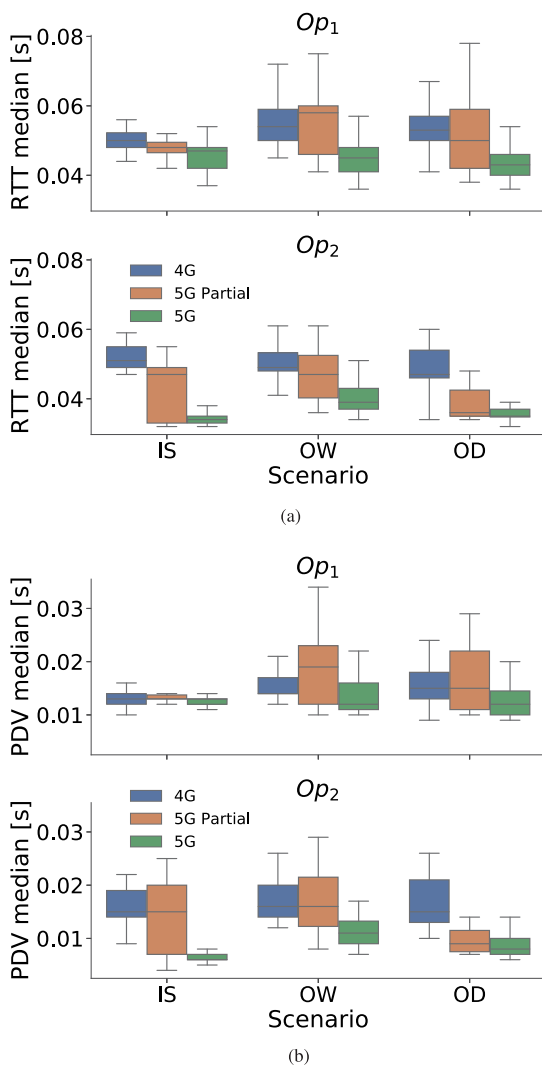


Fig. 13. Real-time online gaming application performance grouped by MNO and scenario. (a) shows the median RTT, (b) shows the PDV. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Fig. 14. (Continued from Fig. 13) Real-time online gaming application performance grouped by MNO and scenario. (c) shows the missed packet percentage, and (d) shows the i -score. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

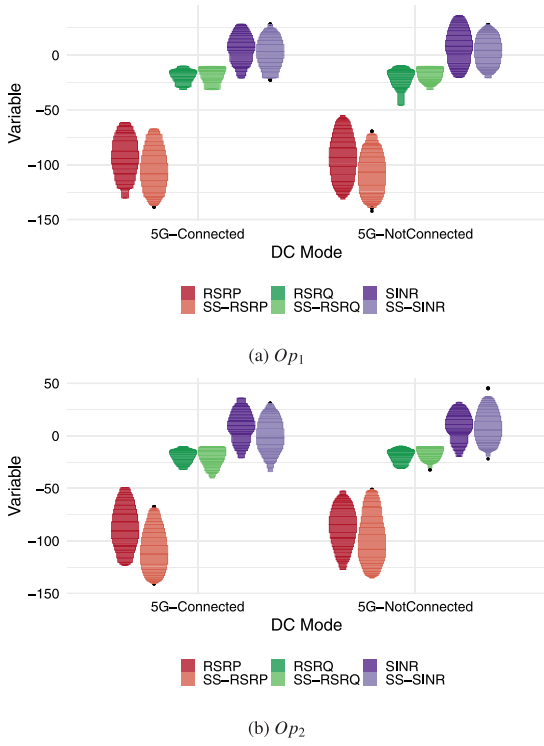


Fig. 15. Letter plots for visualizing the distribution of the LTE and 5G variables, grouped by DC mode and split per MNO; (a) Op_1 and (b) Op_2 . (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

6.2. Exploratory data analysis

We next conduct a series of statistical tests to get a better understanding of the dataset's characteristics but also to determine whether it is feasible to differentiate between the two modes by considering the statistical properties of the radio coverage indicators. Figs. 15(a) and 15(b) depict the distribution of all six variables in a letter-value plot fashion grouped per DC mode for Op_1 and Op_2 , respectively. We observe that across both MNOs, the range of each variable do not varies in a great extent across the two DC modes, a result suggesting that it is rather difficult to identify the DC mode by considering the range of distribution of the radio coverage indicators.

In addition, in Fig. 16 we present the Spearman's rank correlation between each variable of interest using heatmap visualizations. We make the following observations. The correlation between the DC mode and the coverage indicators averages approximately -8×10^{-4} and -0.2 for Op_1 and Op_2 respectively, a result that is in line with Fig. 15. We further observe that the correlation within the groups of LTE variables (RSRP, RSRQ, and SINR) and 5G variables (SS-RSRP, SS-RSRQ, and SS-SINR) is very high for both MNOs, which is expected considering that each variable measures the strength or quality of the cellular network signal from a slightly different perspective. On the other hand, the correlation between the two groups varies significantly across Op_1 and Op_2 (0.406 and 0.008, respectively) which suggests that the MNOs have distinct network configurations in terms of how the 5G RAN is deployed on top of the existing 4G RAN. The above observation motivates our decision to exploit supervised ML for building classifiers towards modeling the DC decision process in 5G NSA networks.

6.3. ML modeling

Our goal is twofold; first, to determine whether it is possible to accurately predict the DC mode using one or more combinations of

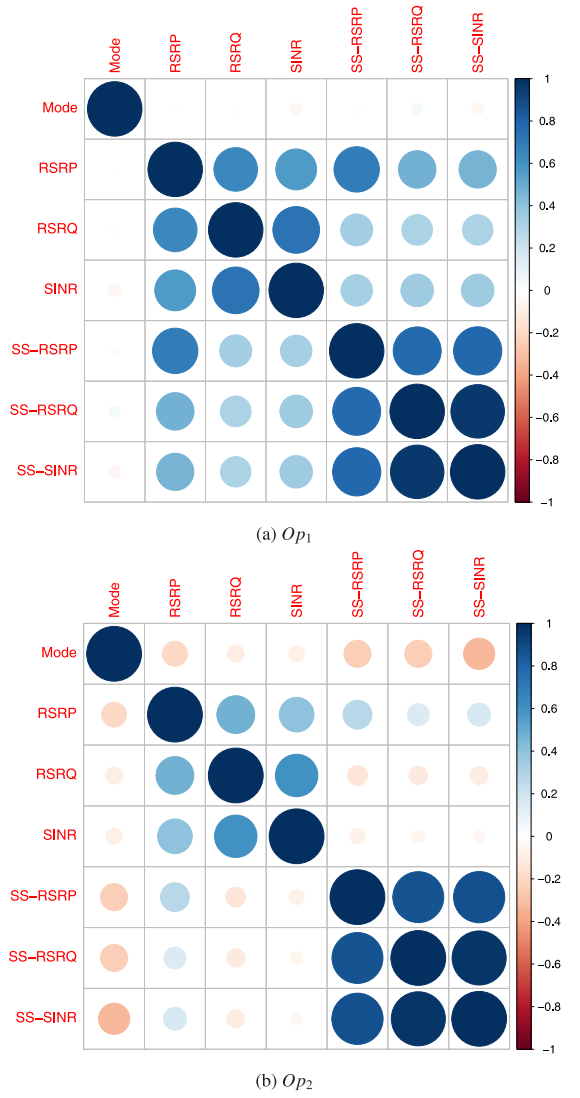


Fig. 16. Heatmaps representing the Spearman's rank correlation between all available variables split per MNO; (a) Op_1 and (b) Op_2 . (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the radio coverage parameters, and second, to pinpoint the importance and impact of the input features to the ML models. We leverage the measurements collected during 5G-enabled sub-campaigns (described in Section 3.2), and define the following two options for the DC Mode.

- *5G-Connected*, for scenarios where the UE was instructed by the network to establish a connection to a 5G PCI on top of an existing connection to one or more 4G PCIs;
- *5G-NotConnected*, for scenarios where the UE was not instructed by the network to establish a connection to a 5G PCI and thus only maintained the existing connection to one or more 4G PCIs.

We further highlight that the second option only includes scenarios where, apart for 4G PCIs, at least one 5G PCI was also detected at the UE's location, thus making DC theoretically possible. We instead exclude scenarios where no 5G PCIs were detected, since in these cases DC cannot be achieved in practice, e.g., due to lack of 5G coverage.

Experimental setup. We test four different ML algorithms: (i) Decision Trees (DTs), which comprise the simplest tree-based ML algorithm, (ii) Support Vector Machines (SVM) [34], which are a set of supervised learning methods that utilize the concept of kernel functions to solve problems involving both linear but also non-linear data, (iii)

Table 4
ML model hyperparameters.

Algorithm	Hyperparameters
DT	–
SVM	Kernel function: Radial
XGB	# trees: 100 Maximum tree depth: 6 Learning rate: 0.3 Tuning of regularization: 0 Column sampling: 1 Minimum leaf weight: 1 Row sampling: 1
RF	# trees: 500 # features: Square root of total features

eXtreme Gradient Boosting (XGB) [35], which is an implementation of distributed Gradient Boosted Decision Trees (GBDT) for improving speed and performance, and (iv) RF [36], which adopts the concept of *ensemble learning* towards improving the performance of DTs.

We evaluate the performance of our models by adopting *repeated k-fold cross-validation* with three repeats and five folds, with the exception of SVM. As SVM models are not tree-based, they require significantly larger run time to train (≈ 200 – 300 times larger in average compared to the other tree-based methods when using the same experimental setup). After testing with different configurations, we found that additional folds and/or repeats had a negligible improvement in SVM model performance, i.e., $<0.5\%$ in terms of accuracy on average. Therefore, considering the tradeoff between performance and run time, we modified the control settings to three folds and zero repeats for SVM. In terms of parameter tuning, we set the number of trees for RF to 500 and the number of features randomly sampled in each individual tree to the square root of the total number of available features used in each model (default), while for SVM, we used the radial kernel function. We summarize the hyperparameters for all four algorithms in Table 4. Last, we split our dataset into training and testing by assigning two-thirds of the data to the former and the remaining one-third to the latter.

6.4. Performance evaluation

The performance evaluation section is divided into three parts; *ML evaluation*, where we evaluate the performance of our ML models using classification accuracy as a metric, *Model explainability*, where we focus on explainability aspects of the RF models, and *Feature Contributions*, where we show the impact of each feature to the classification model.

6.4.1. ML evaluation

Prior to evaluating the performance of our models, we conduct a feature importance analysis via Recursive Feature Elimination (RFE),³ a method that uses a recursive approach to eliminate the weakest feature one at a time. Results show that, for O_{p_1} , SINR is the most important feature, with RSRQ and RSRP being second and third, while for O_{p_2} , RSRP is the most important, followed by SINR and SS-RSRP. For both MNOs, RFE shows that the best performance is achieved by using five (out of six) input features (i.e., all but SS-RSRQ) with an accuracy of 99.07% and 99.34% for O_{p_1} and O_{p_2} , respectively, obtained on the training data.⁴ Even though RFE dictates the use of the above five features for our models, we conduct the analysis using different combinations of the input features pool. By doing so we get a more holistic understanding of the ML models performance when less (or

³ <https://www.rdocumentation.org/packages/caret/versions/6.0-92/topics/rfe>.

⁴ Other feature importance methods, such as the Learning Vector Quantization (LVQ) or the Boruta algorithm (<https://cran.r-project.org/web/packages/Boruta/Boruta.pdf>) provide slightly different results.

more, i.e. all six) features are used while leveraging the results to validate the output of RFE.

Table 5 shows the classification accuracy, defined as the ratio (percentage) of the number of correct predictions to the total number of predictions, obtained for each ML model, algorithm, and MNO. The table is divided into five groups (namely I–V), that differ in terms of the number of input variables used for each model: group I shows results obtained by using a single feature, group II considers two features, group III moves to three features, group IV considers the output features as indicated by RFE, i.e., five features, and finally group V reports results obtained when all the six features are combined. For group II, combinations of input features included both cases where the two features are related to a single technology (labeled as LTE and 5G) or to different technologies (labeled as cross-technology). For each group, we underline the highest accuracy model across each pair of MNO and algorithm, and highlight the highest accuracy model across each MNO using bold text.

Evaluation of ML algorithms. We first compare the performance across the four ML algorithms. We observe that on average, RF outperforms XGB, SVM, and DTs by 6.03%, 22.95%, and 29.25% for O_{p_1} , and 7.02%, 21.55%, and 25.11% for O_{p_2} , respectively. In particular, in none of the models, RF provides inferior performance than the other three algorithms. Among the rest, XGB performs significantly better than SVM, which, in turn, outperforms DTs by a slight margin. As a result, we observe that tree-based models (with the exception of DTs which act as a baseline) provide superior performance, with RF achieving a staggering 99% for O_{p_2} when five parameters are used, a result consistent with the RFE output.

Evaluation across MNOs. Next, we evaluate the models performance across the two MNOs. We find that on average, the difference between O_{p_1} and O_{p_2} for DTs, SVM, and RF is negligible, i.e., 4.67%, 1.93%, and 0.54% in favor of O_{p_2} , while for XGB, the classification accuracy is superior for O_{p_1} by only a slight margin (0.44%). The above result suggests that the two MNOs follow similar policies when it comes to network configuration as the ML models behave similarly across the two datasets.

Evaluation across the number of input variables. Next, we look into the performance of the ML models with respect to the number of input variables. We find that for both MNOs, the average (across all algorithms and models) classification accuracy increases as the number of input variables increases. In particular, for O_{p_1} , the level of increase between groups I and II, II and III, III and IV, and IV and V is 9.26%, 4.15%, 3.72%, and 2.47%, respectively. Likewise, for O_{p_2} , the corresponding difference values are 9.8%, 4.01%, 4.22%, and 1.12%. For group I, both XGB and RF provide the best performance when RSRP is used as the input variable regardless of the MNO. On the other hand, DTs and SVM perform best when SS-RSRP (O_{p_1}) or SS-SINR (O_{p_2}) are used as input. Likewise, for groups II and III, both XGB and RF are superior when LTE variables are used as input (i.e., RSRP and SINR), while the remaining two algorithms achieve their highest accuracy when 5G variables are used. Regardless, the highest performance is offered by RF when five parameters are considered, with an accuracy of 98.6% and 99% for O_{p_1} and O_{p_2} , respectively. The above results highlight the significance of considering the parameters dictated by RFE as input to any of the ML models under study, as each of these parameters, even not equally important, provides valuable information to the model and considerably increase the classification accuracy.

Spatial mapping of the model predictions. We now focus on the RF model that includes all six variables and use a geo-spatial visualization to study how the predictions generated by the model distribute in the space. Fig. 17 shows an example of an OW sub-campaign. We use a different color scheme to represent each predicted DC mode, i.e., orange for *5G-NotConnected*, and purple for *5G-Connected*. We draw black circles to represent the regions where RF predicted the wrong class.

Table 5

Classification accuracy (%) per MNO and ML algorithm. The table is further grouped according to the number of input variables. For each group, the best model across each MNO and ML algorithm is highlighted using underlining, while the best model across each MNO is highlighted using bold text.

Input feature(s)	Op_1				Op_2			
	DT	SVM	XGB	RF	DT	SVM	XGB	RF
I. One input variable								
RSRP	59.3%	60.7%	79%	86.2%	61.1%	60.9%	76.3%	84.5%
RSRQ	54.3%	56%	70.1%	73.4%	55.2%	55.4%	67.8%	71.5%
SINR	55.4%	55.6%	75.9%	81.5%	54.3%	55.1%	71.9%	78.2%
SS-RSRP	<u>59.7%</u>	<u>60.9%</u>	68.9%	70.9%	60.8%	61.4%	69.5%	72%
SS-RSRQ	56.8%	57.6%	64.8%	65.3%	61.4%	61.5%	68.3%	68.4%
SS-SINR	54.1%	56.1%	66.8%	68.3%	<u>65.2%</u>	<u>64.9%</u>	72.1%	72.6%
II. Two input variables								
<i>LTE</i>								
RSRP + RSRQ	59.3%	63.3%	86.5%	94.5%	61.1%	61.4%	85.5%	95.4%
RSRP + SINR	57.2%	65%	<u>89.2%</u>	94.7%	61.1%	62.5%	85.5%	95.5%
RSRQ + SINR	54.9%	60.1%	88.1%	94.4%	54.6%	59%	<u>85.7%</u>	95.1%
<i>5G</i>								
SS-RSRP + SS-RSRQ	59.7%	66.1%	79.6%	87.8%	66.4%	69.2%	79.9%	87.5%
SS-RSRP + SS-SINR	<u>61.7%</u>	66.6%	80.7%	91.9%	66.5%	69.4%	80.6%	93%
SS-RSRQ + SS-SINR	56.5%	<u>68.5%</u>	79%	83.1%	<u>66.7%</u>	<u>73.3%</u>	80.8%	84.5%
<i>Cross-technology</i>								
RSRP + SS-RSRP	59.7%	66.1%	84.3%	94.2%	63.5%	68%	83.4%	95%
RSRQ + SS-RSRQ	54.6%	60.2%	81.3%	88.3%	61.4%	64.5%	80.8%	89.1%
SINR + SS-SINR	56.5%	60.5%	82%	94%	<u>66.7%</u>	68.4%	82.9%	94.5%
III. Three input variables								
RSRP + RSRQ + SINR	57.2%	67.7%	<u>90.8%</u>	96.1%	61.1%	65%	<u>89.5%</u>	96.5%
SS-RSRP + SS-RSRQ + SS-SINR	<u>61.7%</u>	<u>72.7%</u>	85.3%	95.1%	<u>69.3%</u>	<u>76.8%</u>	86.4%	96%
IV. Five input variables (RFE output)								
All (except SS-RSRQ)	57.9%	76.7%	95%	98.6%	67%	78.5%	92.7%	99%
V. Six input variables								
All	64.5%	80.3%	94.8%	98.5%	66.5%	82.3%	94.1%	98.8%




Fig. 17. GPS traces for an example OW sub-campaign. The color of each point represent RF's predicted DC mode, i.e., orange for *5G-NotConnected*, and purple for *5G-Connected*. Black circles highlight areas where the model generates wrong DC mode predictions. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

We observe that, in this example, the errors are localized in two particular regions and are not dispersed across the route. Given that the RF model achieves an accuracy up to 98.8% and thus it can be used to predict the DC mode based on the underlying radio conditions of the nearby cells, we argue that such a mapping can be particularly helpful for MNOs or service providers towards identifying regions where such scenarios of unclear DC mode are likely to take place. Such insights

can be then used for improving deployment and network configuration aspects for future 5G NSA deployments.

6.4.2. Model explainability

The concept of model explainability refers to a set of methods that allow humans to better understand and explain the decisions generated by a ML model. Model explainability not only solves the 'black box' aspect of many ML algorithms, but it can also aid towards debugging and troubleshooting during development.

We focus on explainability aspects of the RF models, considering its superior performance compared to the other three ML algorithms. The most common approach when it comes to explaining the decision-making of a RF model is by plotting its decision boundaries, which are defined as surfaces used to separate data points belonging to a different class. Visualization of decision boundaries is possible by either using a 2-D or a 3-D plane. Figs. 18(a) and 18(b) depict the 2-D decision boundaries with respect to the best performing ML model of group II (Table 5), for Op_1 and Op_2 , respectively, where the *y-axis* is the RSRP while the *x-axis* is the SINR. We use a different color to highlight each decision plane along with the data points belonging to each of the two classes, i.e., *red* for *5G-Connected* and *blue* for *5G-NotConnected*. To reduce visual clutter, we only include test data from a single OD campaign. Note that even though the dataset provides the geographical information (i.e., latitude and longitude) for each measurement point, these are not part of the model's decision making, thus, are not included in the plots.

We observe that for both Op_1 and Op_2 decision boundaries form multiple 'island' areas of variable sizes. These areas do not seem to follow any particular defined patterns, which does not allow for using any sort of threshold values to determine either of the two classes; a

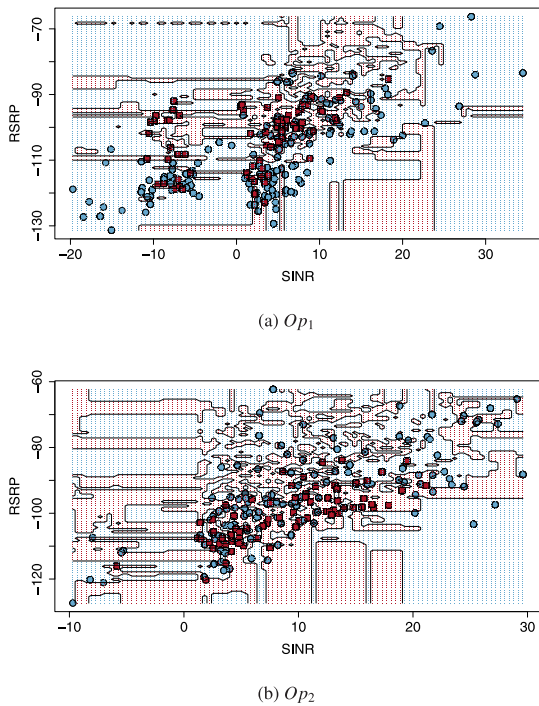


Fig. 18. Decision boundaries of a RF model that uses two independent variables, i.e., RSRP, and SINR, for Op_1 (a) and Op_2 (b). The color indicates the DC mode, i.e., blue for *5G-NotConnected* and red for *5G-Connected*. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

result that further highlights the complexity of the DC decision problem. Across the two MNOs, we observe that decision boundaries present some differences only distinguishable in a high level. For example, for Op_1 , one of the largest *5G-NotConnected* ‘island’ areas is in the lower range of both SINR and RSRP, i.e., $[-20, -12]$ dB and $[-130, -110]$ dBm, respectively, whereas, for Op_2 , data points within the same range have a higher probability of being classified as *5G-Connected*. Even though similar observations can be made for other areas of the figures in a high level, precisely characterizing the operating points of each class is a rather hard task.

6.4.3. Feature contributions

Following the model explainability route, next we measure the impact of each feature to the RF model by looking into aspects related to the concept of feature contributions. We adopt *Forest Floor* [37], a novel methodology for interpreting and visualizing the decisions generated by a RF model. *Forest Floor* relies on feature contributions, which are computed by leveraging information from the tree network of a RF model. RF assigns a probability value to each class and then uses a majority vote mechanism to reach a decision. Therefore, feature contributions are defined as the change of the predicted probability given to each class. Figs. 19 and 20 show the results obtained by *Forest Floor* for two example features (RSRP and SS-RSRP) for Op_1 and Op_2 , respectively. The y-axis shows the additive change of predicted probability for each data sample and for each of the two classes, while the x-axis shows the value range for each feature. The sum of changed probability across the two classes for any data sample adds to zero.

We observe that the change in the predicted probability for RSRP shows the highest variance across the two classes. Between the two MNOs, the probability change (y-axis) for each class follows a different pattern. In particular, for Op_1 , measurement data with reported RSRP between -130 and -105 dBm and -90 and -75 dBm are more likely to be classified as *5G-NotConnected*, while for Op_2 , the respective RSRP

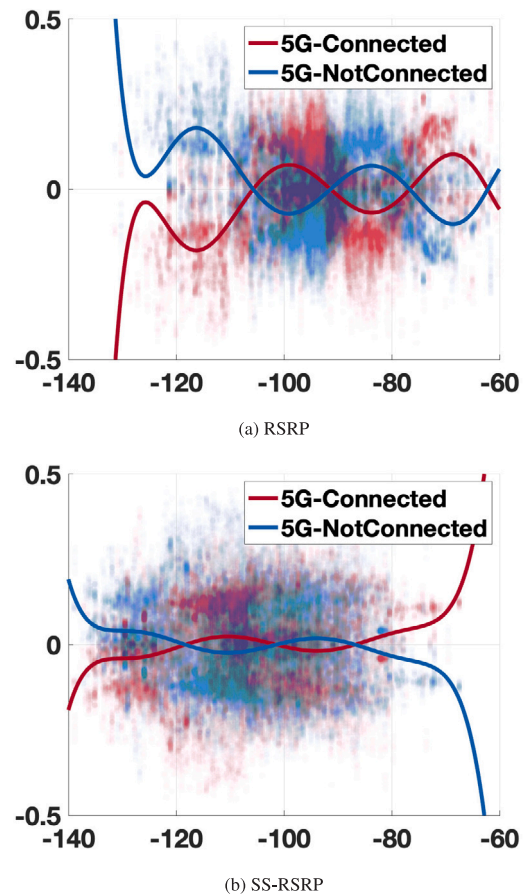


Fig. 19. Feature contributions for two example features (RSRP and SS-RSRP), class (*5G-Connected*, *5G-NotConnected*) and the training data of Op_1 . The y-axis shows the additive change of predicted probability for each data sample and for each of the two classes, while the x-axis shows the value range for each feature. A seventh degree polynomial curve is fit for data shown in each subplot, in order to improve the readability of the results. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

range values are between -90 and -65 dBm. Likewise, other features tend to have different impact across the two MNOs. Overall, we observe that the change of the predicted probability (y-axis) for each class for the SS-RSRP is in average smaller than RSRP, a result that is expected considering that 4G indicators, when used as single predictors in RF models (and XGB), yield the superior performance (see Table 5). However, recall that both figures are visualized using extreme smoothing of degree seven. Therefore, the depicted results should only be used to provide a high-level view of the importance of each feature as the underlying data (blurred in all figures) presents much more complex patterns.

Takeaway Remarks: The above analysis indicates that ML can be efficiently used for predicting the DC mode with an accuracy of up to 99%. Our analyses demonstrate that the combination of five (out of six) radio coverage indicators is instrumental for achieving such a high accuracy. In terms of model explainability, on the one hand, we observe that the analysis of the RF decision boundaries shows that the decision making of the algorithm is extremely complex, thus not allowing to clearly characterize the operating points of each DC Mode class. On the other hand, the results of the *Forest Floor* analysis shows that it is possible to determine the likelihood of predicting a particular class by looking at specific ranges of values for the input parameters. This can be helpful towards better understanding to what extent each parameter contributes to the decision output of each ML algorithm.

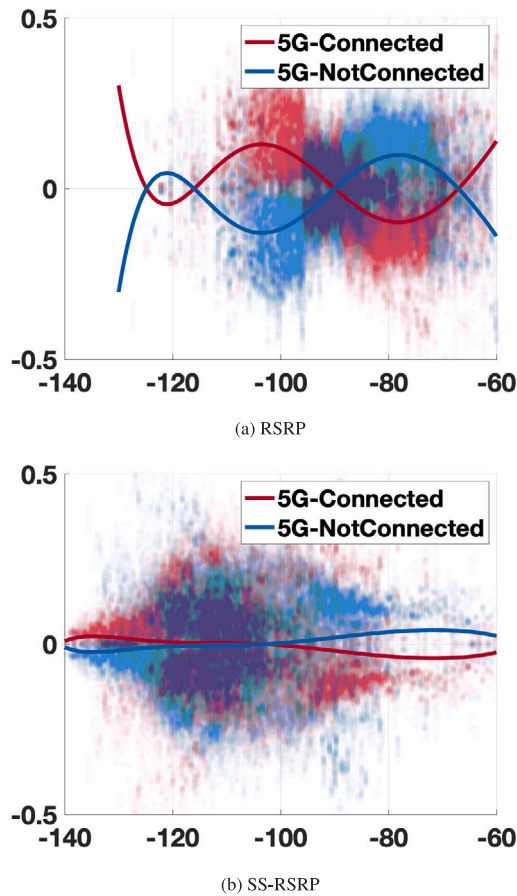


Fig. 20. Feature contributions for two example features (RSRP and SS-RSRP), class (5G-Connected, 5G-NotConnected) and the training data of Op_2 . The y-axis shows the additive change of predicted probability for each data sample and for each of the two classes, while the x-axis shows the value range for each feature. A seventh degree polynomial curve is fit for data shown in each subplot, in order to improve the readability of the results. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

7. Conclusion

This paper presents the first empirical study on network coverage, deployment, and performance of 5G NSA deployments in a European country. We used the collected dataset to study the above aspects for two MNOs in Italy, and make the dataset open-source for further investigation. In addition, we present a ML-based methodology for modeling DC mode using as input the radio coverage characteristics of both 4G and 5G access networks.

Results show that disparate 5G NSA deployment strategies by MNOs lead to different radio coverage performance; among others, we observe that SSB beamforming is effective for improving radio coverage. In terms of QoS/QoE performance, 5G provides significant downlink throughput gains compared to 4G, which is a key requirement for eMBB services. From a URLLC latency/reliability perspective, users may also benefit from using 5G NSA, but coverage holes lead to high performance variability and degradation. Overall, 5G gains over 4G are still unstable particularly in scenarios involving mobility, thus requiring further enhancements and optimization of the NSA architecture and corresponding configurations and procedures. Last, we show that RF can identify the DC mode by achieving an accuracy of up to 99% for Op_2 using as input the radio coverage parameters dictated by RFE, a result that can be used towards building DC connectivity maps on top of existing coverage maps but also to assist MNOs to pinpoint

potential configuration and deployment issues in their 5G NSA network infrastructures.

As a future work, we intend to include more application scenarios, such as video streaming in our active measurements analysis. We also plan to conduct more extensive studies on handovers and DC mode, as well as novel empirical analyses of 5G SA performance.

CRediT authorship contribution statement

Konstantinos Kousias: Writing – review & editing. **Mohammad Rajiullah:** Writing – review & editing. **Giuseppe Caso:** Writing – review & editing. **Ozgu Alay:** Writing – review & editing. **Anna Brunstrom:** Writing – review & editing. **Usman Ali:** Writing – review & editing. **Luca De Nardis:** Writing – review & editing. **Marco Neri:** Writing – review & editing. **Maria-Gabriella Di Benedetto:** Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgments

The work of U. Ali, L. De Nardis and M.-G. Di Benedetto was partially supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, partnership on “Telecommunications of the Future” (PE00000001-program “RESTART”). M. Rajiullah, G. Caso, and A. Brunstrom were supported by the Knowledge Foundation of Sweden. O. Alay and K. Kousias were supported by the EU 6G SNS programme under grant agreement No. 101096452 (IMAGINE-B5G).

References

- [1] G. Liu, Y. Huang, Z. Chen, L. Liu, Q. Wang, N. Li, 5G deployment: Standalone vs. non-standalone from the operator perspective, *IEEE Commun. Mag.* 58 (11) (2020) 83–89.
- [2] K. Kousias, M. Rajiullah, G. Caso, U. Ali, O. Alay, A. Brunstrom, L. De Nardis, M. Neri, M.-G. Di Benedetto, A large-scale dataset of 4G, NB-IoT, and 5G non-standalone network measurements, *IEEE Commun. Mag.* (2023) <http://dx.doi.org/10.1109/MCOM.011.2200707>.
- [3] E. Ali, M. Ismail, R. Nordin, N.F. Abdulah, Beamforming techniques for massive MIMO systems in 5G: Overview, classification, and trends for future research, *Front. Inf. Technol. Electron. Eng.* 18 (6) (2017) 753–772.
- [4] 3GPP, NR; Multi-Connectivity; Overall Description; Stage 2 (Rel-15), *Technical Specification (TS) 37.340*, 2018.
- [5] 3GPP, NR and NG-RAN Overall Description; Stage 2 (Rel-15), *Technical Specification (TS) 38.300*, 2020.
- [6] OpenCellID, 2023, <https://opencellid.org>.
- [7] RTR Netztest, 2023, <https://www.netztest.at>.
- [8] LTE Italy, 2023, <https://lteitaly.it>.
- [9] K. Kousias, M. Rajiullah, G. Caso, O. Alay, A. Brunstrom, L. De Nardis, M. Neri, U. Ali, M.-G. Di Benedetto, Coverage and performance analysis of 5g non-standalone deployments, in: *Proceedings of the 16th ACM Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization*, in: WiNTECH '22, Association for Computing Machinery, New York, NY, USA, ISBN: 9781450395274, 2022, pp. 61–68, <http://dx.doi.org/10.1145/3556564.3558233>.
- [10] 5GPPP, The 5G infrastructure public private partnership, 2023, <https://5g-ppp.eu/>, [Accessed 19-Aug-2023].
- [11] PAWR, Platforms for advanced wireless research - PAWR, 2023, <https://advancedwireless.org/>, [Accessed 19-Aug-2023].
- [12] A. Narayanan, E. Ramadan, J. Carpenter, Q. Liu, Y. Liu, F. Qian, Z.-L. Zhang, A first look at commercial 5G performance on smartphones, in: *Proceedings of the Web Conference 2020*, 2020, pp. 894–905.

- [13] D. Xu, A. Zhou, X. Zhang, G. Wang, X. Liu, C. An, Y. Shi, L. Liu, H. Ma, Understanding operational 5G: A first measurement study on its coverage, performance and energy consumption, in: Proceedings of the ACM SIGCOMM 2020 Conference, 2020, pp. 479–494.
- [14] A. Narayanan, E. Ramadan, R. Mehta, X. Hu, Q. Liu, R.A. Fezeu, U.K. Dayalan, S. Verma, P. Ji, T. Li, et al., Lumos5G: Mapping and predicting commercial mmwave 5G throughput, in: Proceedings of the ACM Internet Measurement Conference 2020, 2020, pp. 176–193.
- [15] A. Narayanan, X. Zhang, R. Zhu, A. Hassan, S. Jin, X. Zhu, X. Zhang, D. Rybkin, Z. Yang, Z.M. Mao, et al., A variegated look at 5G in the wild: Performance, power, and QoE implications, in: Proceedings of the ACM SIGCOMM 2021 Conference, 2021, pp. 610–625.
- [16] A. Hassan, A. Narayanan, A. Zhang, W. Ye, R. Zhu, S. Jin, J. Carpenter, Z.M. Mao, F. Qian, Z. Zhang, Vivisecting mobility management in 5G cellular networks, in: Proceedings of the ACM SIGCOMM 2022 Conference, 2022.
- [17] E. Ramadan, A. Narayanan, U.K. Dayalan, R.A. Fezeu, F. Qian, Z.-L. Zhang, Case for 5G-aware video streaming applications, in: Proceedings of the ACM SIGCOMM 2021 Workshop on 5G Measurements, Modeling, and Use Cases, 5G-MeMU, 2021, pp. 27–34.
- [18] A. Narayanan, E. Ramadan, J. Quant, P. Ji, F. Qian, Z.-L. Zhang, 5G tracker: A crowdsourced platform to enable research using commercial 5G services, in: Proceedings of the ACM SIGCOMM 2020 Poster and Demo Sessions, 2020, pp. 65–67.
- [19] K. Kousias, M. Rajiullah, G. Caso, O. Alay, A. Brunstrom, L. De Nardis, M. Neri, U. Ali, M.-G. Di Benedetto, Implications of handover events in commercial 5g non-standalone deployments in rome, in: Proceedings of the ACM SIGCOMM Workshop on 5G and Beyond Network Measurements, Modeling, and Use Cases, in: 5G-MeMU '22, Association for Computing Machinery, New York, NY, USA, ISBN: 9781450393935, 2022, pp. 22–27, <http://dx.doi.org/10.1145/3538394.3546041>.
- [20] U. Ali, G. Caso, L. De Nardis, K. Kousias, M. Rajiullah, Ö. Alay, M. Neri, A. Brunstrom, M.-G. Di Benedetto, Large-scale dataset for the analysis of outdoor-to-indoor propagation for 5G mid-band operational networks, Data 7 (3) (2022) 34, <http://dx.doi.org/10.3390/data7030034>.
- [21] U. Ali, G. Caso, L. De Nardis, K. Kousias, M. Rajiullah, Ö. Alay, M. Neri, A. Brunstrom, M.-G. Di Benedetto, Data-driven analysis of outdoor-to-indoor propagation for 5G mid-band operational networks, Future Internet 14 (8) (2022) 239, <http://dx.doi.org/10.3390/fi14080239>.
- [22] B. Sliwa, H. Schippers, C. Wietfeld, Machine learning-enabled data rate prediction for 5G NSA vehicle-to-cloud communications, in: Proceedings of the IEEE 5G World Forum, IEEE, 2021, pp. 299–304.
- [23] J. Huang, C.-X. Wang, L. Bai, J. Sun, Y. Yang, J. Li, O. Tirkkonen, M.-T. Zhou, A big data enabled channel model for 5G wireless communication systems, IEEE Trans. Big Data 6 (2) (2018) 211–222.
- [24] Speedtest, 2023, <https://speedtest.net>.
- [25] C. Overturf, How does Speedtest measure my network speeds? 2023, <https://tinyurl.com/4whabuu2>, [Accessed 26-Aug-2023].
- [26] Rohde-Schwarz, QualiPoc android, 2023, <https://tinyurl.com/4w4fyuvf>, [Accessed 26-Aug-2023].
- [27] ETSI, Speech and Multimedia Transmission Quality (STQ); QoS Parameters and Test Scenarios for Assessing Network Capabilities in 5G Performance Measurements, ETSI TR 103 702 V1.1.1 (2020-11), 2020.
- [28] ITU-T, Latency Measurement and Interactivity Scoring Under Real Application Data Traffic Patterns, Recommendation ITU-T G.1051 (2023-03), 2023.
- [29] IETF, A Two-Way Active Measurement Protocol (TWAMP), IETF RFC 5357, 2008, Accessed on January 2024. URL <https://datatracker.ietf.org/doc/html/rfc5357>.
- [30] B. Claise, A. Morton, Packet Delay Variation Applicability Statement, RFC 5481, RFC Editor, 2009, URL <https://rfc-editor.org/rfc/rfc5481.txt>.
- [31] 3GPP, 5G; System Architecture for the 5G System (Rel-15), Technical Specification (TS) 23.501, 2018.
- [32] 3GPP, 5G; NR; Base Station (BS) Radio Transmission and Reception (Rel-15), Technical Specification (TS) 38.300, 2019.
- [33] N.V. Chawla, K.W. Bowyer, L.O. Hall, W.P. Kegelmeyer, SMOTE: synthetic minority over-sampling technique, J. Artificial Intelligence Res. 16 (2002) 321–357.
- [34] M.A. Hearst, S.T. Dumais, E. Osuna, J. Platt, B. Scholkopf, Support vector machines, IEEE Intell. Syst. Appl. 13 (4) (1998) 18–28.
- [35] T. Chen, C. Guestrin, Xgboost: A scalable tree boosting system, in: Proceedings of the 22nd Acm Sigkdd International Conference on Knowledge Discovery and Data Mining, 2016, pp. 785–794.
- [36] L. Breiman, Random forests, Mach. Learn. 45 (2001) 5–32.
- [37] S.H. Welling, H.H.F. Refsgaard, P.B. Brockhoff, L.H. Clemmensen, Forest floor visualizations of random forests, 2016, arXiv preprint [arXiv:1605.09196](https://arxiv.org/abs/1605.09196).



Konstantinos Kousias is a Postdoctoral Fellow at University of Oslo. His research focuses on the empirical modeling and evaluation of mobile and IoT systems using data analytics and AI.



Mohammad Rajiullah is a Senior Lecturer at Karlstad University. His research interests are in the areas of low latency networking, web performance, mobile networks, and IoT.



Giuseppe Caso is a Senior Lecturer at Karlstad University. His research interests include cognitive communications, mobile cellular systems and IoT technologies, and location-based services.



Usman Ali is a Ph.D. student at Sapienza University of Rome. His research interests include rate-splitting multiple access and MIMO.



Ozgu Alay is a Full Professor at University of Oslo, Norway, and a Guest Professor in Karlstad University, Sweden. Her research interests include mobile networks, low latency networking, multipath protocols, and multimodal networked interactions.



Anna Brunstrom is a Full Professor and Research Manager for the Distributed Intelligent Systems and Communications Research Group at Karlstad University, Sweden, and a part-time Distinguished Researcher in the Institute of Software Engineering and Technologies (ITIS) at University of Malaga, Spain. Her research interests include Internet architectures and protocols, low latency communication, multipath communication and performance evaluation of mobile systems.



Marco Neri joined R\&S in 2017 as Application Engineer for Mobile Network Testing. His focus is on 5G and IoT testing worldwide.



Luca De Nardis is an Associate Professor at the DIET Department at Sapienza University of Rome. His research interests focus on indoor positioning, UWB and cognitive communications, and routing.



Maria-Gabriella Di Benedetto is a Full Professor at the DIET Department at Sapienza University of Rome. Her research interests are focused on wireless communications, speech, and signal processing.