

LightGBM for Slice Recognition at 5G PHY and MAC Layers [†]

Rosy Altawil ^{1,2,*}, Lucas Delolme ¹, Vincent Audebert ¹ and Philippe Martins ²

¹ Department of SYSTÈME, Électricité de France Research and Development, 91120 Palaiseau, France; lucas.delolme@edf.fr (L.D.); vincent.audebert@edf.fr (V.A.)

² Information Processing and Communications Laboratory (LTCl), Télécom Paris, 91120 Palaiseau, France; martins@telecom-paris.fr

* Correspondence: rosy.altawil@edf.fr

[†] Presented at the 6th International Conference on Communications, Information, Electronic and Energy Systems, 26–28 November 2025, Ruse, Bulgaria.

Abstract

Slicing functionality makes it possible for an operator to share a 5G physical infrastructure between several virtual networks operated by different institutions. The deployed slices can support a wide range of applications with conflicting QoS targets. The coexistence of these slices on top of a common infrastructure is challenging and remains an open issue. Identifying traffic associated with a given type of slice is required to operate and control network resources in an efficient and secure way. This work proposes new algorithms operating at the physical and MAC layers. The solutions designed identify traffic generated by URLLC and eMBB slices by defining a new LightGBM framework. The algorithms can operate at the base station level in an O-RAN-type architecture. They provide a valuable input to radio resource management and traffic steering procedures.

Keywords: network slicing; traffic classification; PHY; MAC; URLLC; eMBB; LightGBM

1. Introduction

5G networks can be deployed following two different architectures, either Standalone (SA) or Non-Standalone (NSA). 5G NSA depends on a 4G Core network by adding a 5G gNB besides a pre-existing 4G eNB, as can be seen in Figure 1. On the other hand, 5G SA does not depend on a 4G network; instead, it uses its own independent 5G access and core networks.




Figure 1. Difference between SA and NSA 5G.

A main feature introduced with 5G SA is network slicing, which allows the deployment of multiple virtual networks (called slices) on top of a common physical infrastructure [1],



Academic Editors: Teodor Iliev, Ivaylo Stoyanov and Grigor Mihaylov

Published: 20 January 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

each optimised for a specific type of use case called vertical. The two slices studied in this work are eMBB and URLLC.

Enhanced Mobile Broadband (eMBB) is a slice that focuses on users needing high data rates [1]. It is used for data-heavy services, such as online streaming, virtual reality and file transfer. As for Ultra-Reliable Low-Latency Communications (URLLC), it is a slice tailored to critical latency-dependent small-packet communications [2]. It is used for mission-critical applications such as smart grid fault detection, remote healthcare, and self-driving vehicles.

The rest of this paper is organised as follows. Section 2 reviews and compares the existing state-of-the-art articles on the subject, which can be grouped into two families. The first is presented in Section 2.1, which focuses on classification with information access from all layers, and another family is presented in 2.2, where they classify using only information from the low layers. Section 3 introduces the relevant preliminaries, such as the Quality of Service (QoS) framework and the different physical channels relevant to the study. Section 4 dives into the methodology used to develop the presented solution, from the production of the dataset in Section 4.1 to the model and its training in Section 4.2. The results are then discussed in Section 5. And finally, the conclusions are drawn in Section 6.

2. State of the Art

With the introduction of slicing, traffic classification has become crucial to guarantee optimal QoS standards for the different groups of users. The literature shows that traffic classification can be approached following two main axes: either classifying by type of service or by slice type. The related work can be grouped into two families: those that classify the traffic with information access from all layers, and those that do it only using the low layers.

2.1. Classification with Information Access from All Layers

A first approach to slice identification described in [3] relies on multiple Key Performance Indicator (KPI) features such as Guaranteed Bitrate (GBR), roaming and slice type. Several supervised machine learning models were tested and compared, with coarse decision tree giving a perfect precision of 100%. A similar approach can be found in [4] where a Deep Learning (DL) model called DeepSlice is introduced. It is able to identify the network slice even in the case of network failure using high-layer KPIs. The DeepSlice model returned a 91% overall accuracy. Another approach can be found in [5], where an Explainable Machine Learning (XAI) framework was introduced that classifies slices using high-level QoS metrics. This approach allowed for an accuracy of over 99%. However, these models' results rely on information provided from higher layers, which are rarely accessible in practice.

2.2. Classification Using Only Information from Low Layers

Works such as [6,7] focus on an eMBB traffic flow classification by extracting information from the physical (PHY) and MAC layers of the generated traces. Their LightGBM models returned 95% accuracy for [6] (5G NSA) and 89% accuracy for [7] (5G SA). In [8], three machine learning models were tested to classify different applications in an NSA environment having access only to the PHY layer. Their Long Short-Term Memory (LSTM) model returned up to 41% accuracy, the k-Nearest Neighbours (k-NN) model returned up to 50.65% accuracy, and the best results came from the Random Forest (RF) model that returned up to 87.73% accuracy.

On the one hand, the literature shows that traffic service classification has already been greatly investigated. Results have shown excellent accuracy when depending on high

layers, yet showing challenges while relying only on the low layers. On the other hand, traffic slice classification is less discussed, with some works returning near-perfect results when it relies on KPIs obtained from high layers. Approaches using only low layers are yet rarely developed. This shows a clear gap in the literature: the possibility to classify eMBB and URLLC slices with only access to the PHY and MAC layers. This is the contribution of this work, where a LightGBM model is trained and tested on the PHY and MAC layers of a mixed URLLC and eMBB traffic dataset.

Table 1 provides a short classification of the previously stated works.

Table 1. Summary comparison of representative works on 5G traffic classification.

Ref	Focus	Approach	ML Model	Layer Access	Accuracy
[3]	Slice Classification	Supervised ML using high-layer KPI (GBR, roaming, slice ID)	Decision Trees, SVM, k-NN	High Layers	100%
[4]	Slice Classification	DeepSlice model using device/network KPIs (QCI, delay, loss, UE type)	Random Forest, CNN (DLNN)	High Layers	91%
[5]	Slice Classification	Explainable AI (XAI) using high-level QoS (jitter, packet loss)	Random Forest, Decision Trees	High Layers	99%
[6]	Service Classification	LightGBM on PHY-layer records for encrypted flows	LightGBM	PHY	95%
[7]	Service Classification	Refined LightGBM on standalone traffic	LightGBM	PHY/MAC	89%
[8]	Service Classification	Classified different applications in NSA traffic	LSTM, k-NN, Random Forest	PHY	Up to 87.7%
This Work	Slice Classification	LightGBM classification restricted to PHY/MAC traces	LightGBM	PHY/MAC	97%

3. Preliminaries

In 5G, QoS classes are called the 5G QoS Identifier (5QI). Slicing allows for the application of personalised QoS standards for each slice, applying a set of parameters described by the chosen 5QI value. Examples of 5QI values are listed in Table 2 [1].

Table 2. Sample 5QI value list.

5QI	Default Slice	Delay	Error Rate	Priority ¹	Example Services
1	eMBB	100 ms	10^{-2}	20	Conversation Voice
2	eMBB	150 ms	10^{-3}	40	Conversation Video
65	URLLC	75 ms	10^{-2}	7	Mission-Critical Push to Talk Voice (MCPTT-V)
82	URLLC	10 ms	10^{-4}	19	Automation
85	URLLC	5 ms	10^{-5}	21	Electricity Distribution

¹ Lower priority value is higher priority level.

The 5QI value of each slice affects the data provided by the low-level layers of the 5G radio stack (PHY and MAC).

Bidirectional traffic was considered for the experiments. Multiple channels exist in the physical and MAC layers, as defined in [9], and can be seen in Table 3.

Table 3. PHY and MAC layers and their channels.

Layer/Direction	Channel	Traffic Content
PHY/DL	PDSCH PDCCH	User data and UE demodulation information Control information
PHY/UL	PUSCH PUCCH	User data and RAN demodulation information Control information
MAC		Maps logical and physical channels together

Each channel transports either control information or data, both with added indicators for each packet. P_xSCH channels transport user data and demodulation information; in addition, they have indicators such as harq, tb len and prbs; the HARQ process ID; the length of the transmitted block; and the number of PRBs allocated for that packet. P_xCCH channels transport control information regarding the network connection. The MAC layer maps logical and transport channels together to allow traffic of a logical channel to be transmitted over a transport channel. Transport channels are in turn mapped on physical channels by the physical layer.

To distinguish between these slices, a LightGBM (LGBM) model was developed. The function of the LGBM model is explained in Section IV-B. To analyse the results of such a model, the precision, recall, and f1-score are computed from the following values [10]:

- *TP*: Number of True Positives;
- *FP*: Number of False Positives;
- *TN*: Number of True Negatives;
- *FN*: Number of False Negatives.

The precision shows the accuracy of a positive prediction. It gives the ratio between the number of packets that the model has correctly labelled positive, and the total number of packets labelled positive. It is computed using the following formula:

$$P = \frac{TP}{TP + FP} \cdot \quad (1)$$

The recall shows the ability of the model to find positive instances. It shows from all the positive values how many the model correctly identified. The recall is computed using the following formula:

$$R = \frac{TP}{TP + FN} \cdot \quad (2)$$

As for the F1-score, it is a global metric used to evaluate the model in both precision and recall. It is computed using the following formula:

$$F1 = 2 \times \frac{P \times R}{P + R} \cdot \quad (3)$$

Finally, the global accuracy of the model is computed from the following formula:

$$A = \frac{TP + TN}{TP + FP + TN + FN} \cdot \quad (4)$$

4. Method

4.1. Input Dataset Generation

The data used to train and test the model was obtained from a 5G SA network operated in an EDF R&D laboratory. An Amarisoft UE simulator and an Ericsson gNB and core network were used to set up UEs that use either an URLLC or eMBB slice following the 5QI classes.

The radio parameters used in the dataset generation are grouped in Table 4.

Table 4. Radio parameters for dataset generation.

Parameter	Value
Carrier frequency	3891.6 MHz (NR-ARFCN 659424)
System bandwidth	100 MHz
Number of PRBs	273
Subcarrier Spacing	30 KHz

The chosen frequency, 3891.6 MHz, falls in the 5G NR band n78 (3.3–3.8 GHz), which is a widely deployed 5G band due to its availability. The maximum channel bandwidth defined for n78 is 100 MHz as specified in [11], which is what was used in this experiment, as it provides enough resources to study both URLLC and eMBB presence. The number of Physical Resource Blocks (PRB) depends on the bandwidth (100 MHz) and subcarrier spacing (30 kHz), which is why 273 PRBs were chosen to respect the standard NR resource grid.

As for the traffic model chosen, URLLC packet size is very small since it is mission-critical traffic, with a very short inter-arrival time, almost periodic, in addition to some sudden bursts of traffic. As for the eMBB traffic, packets tend to be large with a very variable inter-arrival time, with traffic arriving almost exclusively in bursts. In addition, URLLC traffic requires a lower number of RBs than eMBB traffic.

A total of 1,811,648 packets (1,448,320 URLLC; 363,328 eMBB) were generated; only the PHY and MAC layers were logged into the dataset. The data was later sanitised to fit a dataframe suitable for ML training, as can be seen in Table 5.

Table 5. Simulation traffic profiles used for dataset generation.

Test Number	URLLC Traffic Type	eMBB Traffic Type
1	1 × Ping	None
2	1 × Small UDP	None
3	None	3 × File Download (1 G, 500 M, 100 M)
4	None	1 × Iperf
5	1 × Ping	2 × File Download (500 M, 200 M)
6	1 × Ping	1 × Iperf

To represent URLLC traffic, a mix of ICMP and UDP small-sized packets was used. For eMBB, more demanding applications were used, such as Iperf and large file downloads with cURL. See Table 6 for more details.

Table 6. List of Engineered Features.

Name	Description
inter_packet_time_ns	Time elapsed since previous packet arrival (in ns)
time_to_next_packet_ns	Time until next packet arrival (in ns)
harq_rolling_mean	Mean HARQ process ID
harq_rolling_std	Std. deviation of HARQ process ID
prbs_rolling_mean	Mean number of PRBs allocated
prbs_rolling_std	Std. deviation of PRBs allocated
tb_len_rolling_mean	Mean transport block length (bits)
tb_len_rolling_std	Std. deviation of transport block length (bits)
mod_rolling_mean	Mean modulation order
mod_rolling_std	Std. deviation of modulation order
rv_idx_rolling_mean	Mean redundancy version index
rv_idx_rolling_std	Std. deviation of redundancy version index
cr_rolling_mean	Mean code rate
cr_rolling_std	Std. deviation of code rate
retx_rolling_mean	Mean retransmission count
retx_rolling_std	Std. deviation of retransmission count
crc_ok_rolling_mean	Mean CRC success indicator
crc_ok_rolling_std	Std. deviation of CRC success indicator
snr_rolling_mean	Mean signal-to-noise ratio
snr_rolling_std	Std. deviation of signal-to-noise ratio
mcs_rolling_mean	Mean modulation and coding scheme index
mcs_rolling_std	Std. deviation of modulation and coding scheme index
format_rolling_mean	Mean PUCCH format indicator
format_rolling_std	Std. deviation of PUCCH format indicator
total_len_rolling_mean	Mean MAC packet length (bytes)
total_len_rolling_std	Std. deviation of MAC packet length (bytes)

4.2. Model and Training

After generating a sample dataset, the goal was to build a simple machine learning classifier to separate packets of a mixed data stream into URLLC and eMBB. An LGBM model was chosen for its great compromise between training time, complexity and accuracy. The LGBM can achieve such good accuracy in a small training time thanks to optimised built-in mechanisms.

First is its leafsplit strategy, as can be seen in Figure 2. This allows for better accuracy by pruning paths that have a great chance of becoming dead ends.




Figure 2. Chart showing LightGBM's leafsplit strategy.

Next is the model's Exclusive Feature Bundling (EFB), which mainly groups similar features together. It allows the model to learn more easily. In addition to that, in [4,5], similar work was performed to classify eMBB traffic using LightGBM, returning a satisfactory accuracy of 89% and 95%. Finally, a Gradient Boosting Decision Tree (GBDT) was used with a bagging data sample strategy and a learning rate of 0.01. Since the model is used to classify between URLLC or eMBB traffic, the objective is binary.

Training of the model started by feeding it the complete, unmodified, low-level traces. However, it soon got into trouble differentiating small-sized eMBB packets from URLLC traffic. This is why some contextual packet features were engineered like "Time to next packet" and "Rolling packet size mean". These features were built working in 5 ms windows around each packet to gather a sufficient amount of data while keeping data size and model I/O latency in a reasonable range (for possible real-time applications). The full list of these engineered features is available in Table 6.

To reduce confusion between small eMBB and URLLC packets, the engineered features were computed on a rolling window of 5 ms. That approach allowed the model to study not only the instantaneous traffic properties but also the temporal patterns over a window of 5 ms. In the following section, the impact of these features on the LGBM model is shown, and the results reached after the testing of the model are explained and analysed.

5. Results

The model was tested using the designed dataset, and the features were ranked by how important they were to the classification process, as can be seen in Figure 3.

The most prominent features are `tb_len` and `total_len`. They represent the length of transmitted blocks. Since URLLC packets are nearly always small and eMBB packets are mostly large, the size is the first way to distinguish the slice.

The second group of important features is the rolling statistics (`tb_len_rolling_std` and `total_len_rolling_std`) since they reveal traffic patterns over time. These indicators show the variability of the feature in the time window, helping detect bursts or stability in the traffic by showing how packet size changes in a sequence of packets. Since a packet's size can overlap between eMBB and URLLC (small eMBB packets, such as a ping), the rolling statistics help differentiate between the two. URLLC traffic, being mostly stable, would

have a rolling standard deviation around zero; as for eMBB traffic and its large packets and variability in size, the standard deviation tends to be high.




Figure 3. Bar graph of the LightGBM model key features by importance.

Another important feature is `harq_rolling_std`. It is computed using the IDs of active hybrid HARQ processes in the rolling window. A high variability of this indicator can identify high throughput as multiple HARQ processes are running concurrently. In contrast, a small variability can show very low data rates.

By referring to these features, on top of the other less important features, the LGBM model was able to classify the traffic between URLLC and eMBB with a very high accuracy, as can be seen in the confusion matrix in Figure 4.




Figure 4. Confusion matrix of the model.

To use the formulas mentioned in Section 3, URLLC is considered as “Positive” and eMBB is considered as “Negative”. The model gave true URLLC $T_U = 1.4 \times 10^6$, false URLLC $F_U = 27,529$, true eMBB $T_e = 335,799$ and false eMBB $F_e = 23,873$, which leads to an accuracy of $A = 0.9712$. An accuracy of 97.12% is higher than the accuracies of the previous works discussed in the State of the Art Section, such as 95% in [6] and 89% in [7], where eMBB traffic was classified with access to only the MAC and PHY layers. The same goes for [8], where the highest accuracy reached was 87.73% to characterise the services in NSA traffic. As for [3], where an accuracy of 100% is reached, it is important to note that the work was performed with access to all the layers, especially the KPI features such as the slice type.

The classification report is grouped in Table 7, in which precision, recall and F1-score of both slices' prediction are included.

Table 7. Classification report of the LightGBM traffic classifier.

Class	Precision	Recall	F1-Score
URLLC	0.98	0.98	0.98
eMBB	0.93	0.92	0.93

For the URLLC slice's results, the precision is at 0.98, which shows that when the model predicts that a packet is URLLC, it is correct 98% of the time, meaning only 2% of eMBB packets are falsely classified as URLLC. The recall at 0.98 shows that from all the URLLC packets it received, the model correctly labelled 98% of them; only 2% were mislabeled as eMBB. These values led to an F1-score of 0.98, showing that the URLLC traffic was easy to identify due to its stability and small-sized packets.

As for the eMBB slice's results, the precision is at 0.93, meaning that when the model predicts that a packet is eMBB, it is right 93% of the time; 7% of URLLC packets are falsely labelled as eMBB. As for the recall, it is at 0.92, meaning that about 8% of eMBB packets are falsely flagged as URLLC. These values led to an F1-score of 0.93, which is still a high score but lower than the one reached for URLLC traffic. The F1-score reflects the variability of eMBB traffic, given that it can in some cases mimic URLLC traffic.

6. Conclusions

This article proposed a LightGBM machine learning model that can distinguish between URLLC and eMBB packets, while having access only to the MAC and PHY layers. Since real-life data was not obtainable, the dataset was generated using realistic radio parameters.

The LGBM model achieved a 97% accuracy, with a higher accuracy for URLLC classification than eMBB, with features such as the size and the rolling statistics being dominant indicators. This model outperforms similar studies on the state of the art of low-level layers and shows that a very high accuracy can be reached without access to high layers.

It is important to mention that this work was performed on an experimental dataset. The same code can be used in the future on an operational infrastructure. Moreover, this model only focuses on the two slices URLLC and eMBB. In the future, it could be developed to include more slices, such as mMTC.

Author Contributions: Conceptualization, R.A., P.M. and V.A.; methodology, R.A. and L.D.; software, L.D., R.A.; validation, P.M. and V.A.; formal analysis, R.A.; investigation, R.A. and L.D.; data curation, L.D.; writing—original draft preparation, R.A. and L.D.; writing—review and editing, P.M. and V.A.; visualisation, R.A.; supervision, P.M. and V.A.; project administration, P.M. and V.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The simulation traces and source code used in this work are publicly accessible at <https://github.com/superdarki/5g-sa-traffic-classification> (accessed on 16 January 2026).

Acknowledgments: This work benefited from the use of Ericsson's gNB and core network at EDF Chatou R&D, which were essential for the generation of the dataset.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript

SA	Standalone
NSA	Non-Standalone
eMBB	Enhanced Mobile Broadband
URLLC	Ultra-Reliable Low-Latency Communications
QoS	Quality of Service
KPI	Key Performance Indicator
GBR	Guaranteed Bitrate
DL	Deep Learning
PHY	Physical
LSTM	Long Short-Term Memory
k-NN	k-Nearest Neighbours
RF	Random Forest
5QI	5G QoS Identifier
PRB	Physical Resource Block
EFB	Exclusive Feature Bundling
GBDT	Gradient Boosting Decision Tree

References

1. 3rd Generation Partnership Project (3GPP). Technical Specification Group Services and System Aspects; System Architecture for the 5G System (5GS); Stage 2 (Release 17). Technical Specification TS 123 501 V17.14.0. ETSI, Sept. 2024. Available online: https://www.etsi.org/deliver/etsi_ts/123500_123599/123501/17.14.00_60/ts_123501v171400p.pdf (accessed on 12 September 2025).
2. 3rd Generation Partnership Project (3GPP). 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Release 15 Description; Summary of Rel-15 Work Items. Technical Report TR 21.915 V15.0.0. Release 15. 3GPP, Sept. 2019. Available online: https://www.3gpp.org/ftp/Specs/archive/21_series/21.915/21915-f00.zip (accessed on 12 September 2025).
3. Mohammedali, N.A.; Kanakis, T.; Al-Sherbaz, A.; Agyeman, M.O. Traffic Classification Using Deep Learning Approach for End-to-End Slice Management in 5G/B5G. In Proceedings of the 2022 13th International Conference on Information and Communication Technology Convergence (ICTC), Jeju Island, Republic of Korea, 19–21 October 2022; pp. 357–362. [CrossRef]
4. Anurag, T.; Walunj, V.; Beard, C. DeepSlice: A Deep Learning Approach Towards an Efficient and Reliable Network Slicing in 5G Networks. In Proceedings of the 2019 IEEE 10th Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON), New York, NY, USA, 10–12 October 2019. [CrossRef]
5. Ray, D. Network Slice Recognition with Explainable Machine Learning. *J. Adv. Artif. Intell.* **2024**, *1*, 36–50. Available online: <https://jaaionline.org/archives/volume1/number7/ray-jaai202436.pdf> (accessed on 16 January 2026).
6. Xiao, F.; Martins, P.; Lu, J. Real-Time Traffic Classification for 5G NSA Encrypted Data Flows with Physical Channel Records. In Proceedings of the 2023 IEEE 98th Vehicular Technology Conference (VTC2023-Fall), Hong Kong, China, 10–13 October 2023; pp. 1–6. [CrossRef]
7. Aphayavong, B.; Fei, X.; Lu, J.; Martins, P. Optimizing LGBM for Multi-Classification of 5G SA Traffic. In Proceedings of the 2024 IEEE 100th Vehicular Technology Conference (VTC2024-Fall), Washington, DC, USA, 7–10 October 2024; pp. 1–7.
8. Marañón, J.A.T.; Duvignau, R. Classifying 5G Encrypted Packet Traces. In Proceedings of the 2024 International Conference on Electrical, Communication and Computer Engineering (ICECCE), Kuala Lumpur, Malaysia, 30–31 October 2024; pp. 1–6. [CrossRef]
9. 3rd Generation Partnership Project (3GPP). NR; NR and NG-RAN Overall Description; Stage-2; Technical Specification TS 38.300 V17.3.0; ETSI: Sophia Antipolis, France, 2023.
10. Grandini, M.; Bagli, E.; Visani, G. Metrics for Multi-Class Classification: An Overview. *arXiv* **2020**, arXiv:2008.05756. [CrossRef]
11. 3rd Generation Partnership Project (3GPP). NR; User Equipment (UE) Radio Transmission and Reception; Part 1: Range 1 Standalone; Technical Specification TS 38.101-1 V17.10.0; ETSI: Sophia Antipolis, France, 2023.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.