



# AI-driven orchestration at scale: Estimating service metrics on national-wide testbeds

Rodrigo Moreira<sup>a</sup> , Rafael Pasquini<sup>b</sup> , Joberto S.B. Martins<sup>c</sup> , Tereza C. Carvalho<sup>d</sup> , Flávio de Oliveira Silva<sup>e</sup>

<sup>a</sup> Federal University of Viçosa (UFV), Rio Paranaíba, Minas Gerais, Brazil

<sup>b</sup> Federal University of Uberlândia (UFU), Uberlândia, Brazil

<sup>c</sup> Salvador University (UNIFACS), Salvador, Bahia, Brazil

<sup>d</sup> University of São Paulo (USP), São Paulo, Brazil

<sup>e</sup> University of Minho (UMinho), Braga, Portugal

## ARTICLE INFO

### Keywords:

Network slicing  
Deep Neural Networks  
Machine learning  
Service-level agreement  
Distributed database

## ABSTRACT

Network Slicing (NS) realization requires AI-native orchestration architectures to efficiently and intelligently handle heterogeneous user requirements. To achieve this, network slicing is evolving towards a more user-centric digital transformation, focusing on architectures that incorporate native intelligence to enable self-managed connectivity in an integrated and isolated manner. However, these initiatives face the challenge of validating their results in production environments, particularly those utilizing ML-enabled orchestration, as they are often tested in local networks or laboratory simulations. This paper proposes a large-scale validation method using a network slicing prediction model to forecast latency using Deep Neural Networks (DNNs) and basic ML algorithms embedded within an NS architecture evaluated in real large-scale production testbeds. It measures and compares the performance of different DNNs and ML algorithms, considering a distributed database application deployed as a network slice over two large-scale production testbeds. The investigation highlights how AI-based prediction models can enhance network slicing orchestration architectures and presents a seamless, production-ready validation method as an alternative to fully controlled simulations or laboratory setups.

## 1. Introduction

Modern applications require challenging behaviors from physical networks to satisfy stringent requirements such as ultra-reliability, low latency, and high throughput [1]. In addition to these quantifiable network requirements, it is necessary to incorporate seamless, intelligent, and pervasive network capabilities to satisfy user demands [2,3]. Although network management, control planes, and data planes have evolved to address this issue, challenges remain and require further large-scale evaluation.

Many approaches, technologies, and methods have been developed to build user-oriented network architectures that provide connectivity in an isolated and personalized manner [4]. One key technological enabler of this vision is network slicing, which establishes network connectivity on top of physical infrastructure while ensuring isolation, end-to-end connectivity, and application-driven requirements, with dedicated control and data planes [5]. With this service-tailoring capability,

Machine Learning (ML) effectively addresses various management and orchestration challenges, thereby enabling intelligent and real-time insights for service provider managers. A major advantage of intelligent network orchestration is the ability of Artificial Intelligence (AI) to ensure and evaluate the service quality and user experience in network slicing [6].

AI techniques, such as reinforcement learning, supervised learning, and unsupervised learning, have been effectively integrated with network orchestrators to mitigate cybersecurity threats, enable intelligent resource allocation, and ensure Service-Level Agreement (SLA) assurance for network slicing [7–9]. In network slicing SLA assurance and measurement, computational intelligence techniques have been progressively incorporated into the building blocks of orchestration architectures. However, these architectures are neither natively secure nor inherently intelligent, often leading to context-specific solutions [10]. Consequently, they remain dependent on third-party vendors, raising concerns regarding the security and privacy of network slicing services.

\* Corresponding author.

E-mail addresses: [rodrigo@ufv.br](mailto:rodrigo@ufv.br) (R. Moreira), [rafael.pasquini@ufu.br](mailto:rafael.pasquini@ufu.br) (R. Pasquini), [joberto.martins@animaeducacao.com.br](mailto:joberto.martins@animaeducacao.com.br) (J.S.B. Martins), [terezacarvalho@usp.br](mailto:terezacarvalho@usp.br) (T.C. Carvalho), [flavio@di.uminho.pt](mailto:flavio@di.uminho.pt) (F. de Oliveira Silva).

<https://doi.org/10.1016/j.future.2025.107971>

Received 30 September 2024; Received in revised form 11 February 2025; Accepted 9 June 2025

Available online 20 June 2025

0167-739X/© 2025 Elsevier B.V. All rights reserved, including those for text and data mining, AI training, and similar technologies.

A key challenge in incorporating intelligence into network slicing architectures for the effective utilization of AI pipelines is ensuring data quality and granularity [11] while preserving user privacy and avoiding the exposure of network slicing application details. Our approach addresses this by leveraging generic infrastructure metrics appropriately combined to enable SLA fitting and forecasting across diverse network slicing applications. Although demonstrated through SLA prediction for Cassandra applications, the proposed methodology is designed with principles and techniques that can be generalized to various slicing applications. Evaluating the effectiveness of this intelligent orchestration mechanism requires analyzing network environments under conditions comparable to those of a production-ready network [12].

This paper investigated the performance of our intelligent native orchestration architecture in production-ready network scenarios under realistic conditions. Hence, we developed and enhanced an orchestration framework for network slicing [13]. We evaluated the capability of the newer Predictor Module within the “SF12 AI Management” building block to forecast latency from a distributed database (Cassandra) deployed as a network slice, operating on nationwide testbeds FIBRE-NG [14] and Fabric [15]. DNNs and basic ML algorithms were employed to predict Cassandra latency within our dataset generation framework.

This paper presents several key contributions: (1) an empirical analysis of DNNs and basic ML algorithms for network slicing latency forecasting on large-scale testbeds, (2) an in-depth investigation of hyperparameter tuning for DNNs, (3) the development of a realistic network application dataset with a comprehensive workload, and (4) a detailed deployment workflow for network slicing services on nationwide testbeds.

The structure of the remainder of this paper is as follows. In Section 2, we contextualize our work within a broader scope, highlighting the unique contributions of this research. The proposed method is presented in detail in Section 3, followed by a description of the experimental setup in Section 4. Section 5 discusses the results, and Section 6 presents the concluding remarks and future directions.

## 2. Related work

In the literature, there are different uses of ML for predicting Quality of experience (QoE) or Quality of Service (QoS), both for applications in the context of mobile networks and generic applications on network slices [16]. Some approaches use machine learning algorithms pre-trained in simulation environments, validate how these algorithms behave, and compute performance metrics [17]. Other approaches have been proposed for network simulation environments, leaving aside the validation of these methods in real network testbeds [18].

Pasquini and Stadler [19] combined three metrics to estimate application QoE while running different applications in an OpenFlow network. They combined the operating system metrics, network flows in switch tables, and application metrics to train machine learning. The experiments considered two applications: Video on Demand (VoD) and Voldemort-distributed databases. They simulated two traffic behaviors to assess the suitability of QoE prediction by considering the observed metrics. Our approach further explores and validates whether ML behaves properly in a real network testbed.

In network slicing orchestration matter, Cui et al. [20] proposed reinforcement learning combined with the Long Short-Term Memory (LSTM) algorithm to guarantee QoS for vehicle-to-vehicle network slicing. The search space of Deep Reinforcement Learning (DRL) is a 5G antenna, where the action relies on the number of resource blocks. Our approach estimates QoE in a real network testbed.

Nougnanke et al. [21] proposed and evaluated an ML-based approach for modeling and predicting network traffic under different workloads. They can estimate latency under different network conditions using ML algorithms. They used a mininet and a simulated environment to validate their findings using three different datasets.

Our approach involves creating a dataset to validate the estimation of some network metrics; however, our approach focuses on estimating the QoE in a real sliced network.

Ge et al. [22] proposed an Graph Neural Networks (GNN)-based end-to-end delay estimator for Software-Defined Networking (SDN) environments. They used a real dataset from the GEANT network to fit their model. In addition, as a proof of concept, they used the OMNet++ simulation environment to assess the prediction accuracy of the model for QoS assurance. A similar study using the same method as described above was differentiated using the Abline dataset [23].

Laiche et al. [24] introduced a multifactor influence in video QoE. They used machine learning to predict QoE assurance in an emulated network scenario. They used and compared the performance of the K-Nearest Neighbors (KNN), Random Forest (RF), and Decision Tree (DT) algorithms to predict user experience by considering popularity and user engagement on a well-known streaming web platform.

Abdelwahed et al. [25] proposed and evaluated an ML-based approach for estimating Web QoE using different context metrics. The metrics considered were network, browsing, and web user engagement. Using RF, DT, and KNN, they estimated Mean Opinion Score (MOS) by considering user-side metrics. Similarly, our approach aims to predict the application QoE by considering different metrics.

Regarding QoS estimation in network slicing using ML algorithms, Khan et al. [26] applied and evaluated Deep Neural Network (DNN) methods to select better network slicing according to connectivity services (mMTC, URLLC, and eMBB) requirements. Their experiments assessed how the ML algorithm handled accurate slice assignment, slice-load balancing, and slice-failure scenarios. Instead of using a simulated network, our approach evaluates the effectiveness of acML in a real network to predict the applications QoE.

For 5G mobile network slicing, Thantharate and Beard [27] proposed and evaluated a transfer learning method for network slicing QoS. Their proposed rationale is based on training local models for different slice network requirements in the source domain. The model can predict the resources in the target domain to satisfy slicing quality agreement. Our approach further predicts the application of SLA conformance for an online network slice application in a real network testbed. Similarly, N. P. Tran et al. [28] proposed and evaluated an approach to estimate end-to-end throughput in 5G and B5G using LSTM.

Yu et al. [29] idealized a linear regression algorithm combined with Reinforcement Learning (RL) to predict slice mobility while ensuring QoS for the application, minimizing costs, and maximizing revenues and profits. They used different statistics, such as user demand for CPU, RAM, and slicing resources to train the ML algorithm offline. Our approach focuses on predicting application QoE through time-series regression, which enables multi-step predictions over time.

Yang et al. [30] bring a method for predicting QoS for Virtual Reality applications in a network slicing using different machine learning algorithms. Their solutions can predict the latency, bandwidth, and different video codecs for different slicing models. Our approach aims to predict the Cassandra application SLA in a real network testbed by using ML algorithms.

Dangi and Lalwani [31] proposes a hybrid deep learning model for efficient network slicing in 5G networks. The model combines Harris Hawks Optimization (HHO) with Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) networks to optimize hyperparameters and classify network slices. The methodology involves three phases: loading the dataset, optimizing it with HHO, and classifying slices using the hybrid model. The results demonstrate that the proposed model outperforms existing methods in predicting appropriate network slices and offers improved service quality and efficiency in 5G network slicing.

Baktir et al. [32] proposes a Mixed Integer Programming (MIP) model and a heuristic algorithm called NESECS to optimize network slicing for various service types with different performance requirements. The MIP model aims to provide optimal solutions for small network instances, whereas NESECS is designed to handle larger instances

**Table 1**  
Related works and their properties.

Paper	Experimental environment	Evaluation metrics	Dataset	Enabling technologies	Applications
[20]	Simulated network	Convergence time and Mobility velocity	Own	LSTM-DDPG	V2V communication
[21]	Emulated network	Normalized Mean Absolute Error (NMAE), and Prediction Time	Telecom Italia Big Data Challenge dataset [33]	SDN, P4, NS3, and ML Algorithms.	Network connectivity
[22]	Emulated network	RMSE, MAE, and Weighted Mean Absolute Percentage Error (WMAPE).	GEANT	GNN, MLR, XGBoost, and RF	Network connectivity
[23]	Simulated network	Packet Delay, Mean Squared Error (MSE), RMSE, MAE, and WMAPE.	Abilene network	GNN, RF, and NN	Network connectivity
[19]	Local network	NMAE, and Training time	Own	OpenFlow, and ML Algorithms	VoD, and KV
[24]	Simulated network	RMSE, Correlation, and Outlier Ratio	Own	KNN, Decision Three (DT)	VoD
[25]	Local network	Accuracy, Correlation, and MOS rating.	Own	KNN, RF, and DT	Web
[26]	Simulated network	Accuracy, Recall, Precision, and F-1 score	DeepSlice, and Secure5G	Convolutional Neural Network (CNN), LSTM and Python-based ML frameworks.	5G verticals
[27]	Simulated network	MSE, correlation	Own	Matlab deep Learning toolbox.	5G verticals
[29]	Simulated network	RL-Convergence time	Own	NetworkX	5G verticals
[28]	Simulated network	Mean Absolute Percentage Error (MAPE), packet loss, and delay	Own	LSTM, and Traditional ML models	5G verticals
[30]	Simulated network	Accuracy	Own	Linear regression methods	Virtual reality
[31]	Simulated network	Accuracy, Precision, Recall, and F1-Score	Unicauca IP Flow, and 5G Network Slicing	HHO, CNN, and LSTM	Network slicing
[32]	Simulated network	SLA Violations, End-to-End Delay, and Resource Utilization	Custom Simulated data	SDN, Virtualization, and Edge computing	Network slicing, and Edge computing optimization
<b>Our approach</b>	<b>Production-ready network</b>	<b>MAE, MAPE, and MSE</b>	<b>Own</b>	<b>DNNs, and Basic ML models</b>	<b>Distributed Database on a sliced Testbed</b>

efficiently. This study evaluates the performance of these solutions through extensive experiments, demonstrating that the proposed methods can effectively manage network resources, reduce SLA violations, and improve the overall network performance by ensuring logical isolation and resource reservation for different service types.

Approaches in the literature close to this proposal are summarized in Table 1. In this table, we have an “Experimental Environment” column related to which network environment the ML algorithm makes QoS or QoE predictions. Many approaches have used simulated environments or local networks. The “Evaluation Metrics” column aims to show which metrics (MAPE, MSE, RMSE, and others) the authors most often take into account to validate their proposals. The “Dataset” column indicates the dataset that each approach considers in the model training phase. Some approaches have created their own datasets, whereas others use datasets created by third parties. The “Enabling Technologies” column aims to identify which methods and technologies for both machine learning and simulating network environments are state-of-the-art. Many approaches still use classic ML approaches such

as RF, DT, and KNN. Finally, the “Application” column reports the applications for which the prediction engines estimate QoS and QoE. Some approaches focus on 5G verticals, while others focus on specific applications.

### 3. Proposed method

Deploying network slices across multiple domains still requires advanced management and orchestration technologies capable of influencing the underlying network, particularly when dealing with heterogeneous devices. Existing tools and techniques for implementing dynamic and elastic slices remain inadequately managed, presenting opportunities for improvement, especially with AI as a foundational enabler of such architectures. In this context, we previously proposed the Slicing Future Internet Infrastructures (SFI2) reference architecture to manage and orchestrate AI-native network slices while integrating diverse testbeds [13].

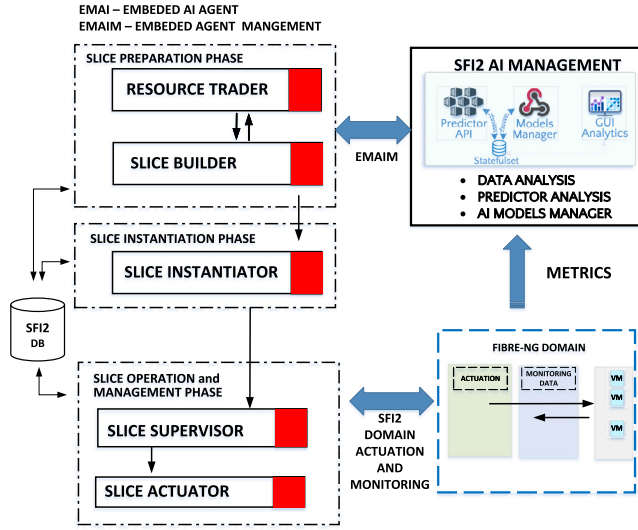


Fig. 1. SFI2 artificial intelligence agents management within FIBRE-NG domain slices.

### 3.1. SFI2 slicing architecture

The SFI2 architectural approach with native artificial intelligence embedded agents is illustrated in Fig. 1. AI-native embedded agents in the architecture target the preparation, commissioning, operation, and decommissioning phases of the network slicing life-cycle.

Enhancements promoted by AI-embedded agents include predicting resource and performance parameters using distinct AI techniques, such as combinatorial optimization, reinforcement learning, and neural networks. Notably, a potential focus of AI-embedded agents is the orchestration process involving different steps and actions applied to instantiate and dynamically reconfigure slices according to user requirements.

Furthermore, the SFI2 architecture aims to operate over heterogeneous infrastructure following the concept of Machine Learning as a Service (MLaaS). To achieve this, the SFI2 AI management module collects metrics from the target domains. It interacts with the embedded agents in the functional blocks to manage the learning model and the other required parameters. The SFI2 AI management module manages learning agents throughout all infrastructure components to support the training, prediction, or decision tasks.

### 3.2. Problem setting and method

This paper proposes the evaluation of network slicing latency forecasting in a sliced SFI2-conform large-scale production-ready testbed (FIBRE-NG and Fabric). The focus is on the SFI2 AI management functional block, which natively and intelligently orchestrates slices to estimate the SLA compliance of an application running on a network slice.

The proposed method is shown in Fig. 2, highlighting the functional block of the orchestrator architecture and its interaction with the network slice lifecycle.

Thus, adopting the premise that the slice is implemented by the SFI2 Orchestrator [13], the procedures for online forecasting of QoE begin by considering different metrics in different testbeds. The step ① of the method refers to the collection and aggregation procedure of different  $X$  metrics from the application, the network slice, and the underlying computational infrastructure. Application metrics  $X_{application}$  refer to the statistics provided by applications at different execution stages. The  $X_{cluster}$  infrastructure metrics refer to the computational resource consumption metrics demanded by the network slices from the underlying hardware. Network metrics  $X_{network}$  refer to the network statistics of the network slice. These metrics are aggregated by data

cleaning and standardization algorithms in the computational resource, resulting in a dataset  $X = X_{application} \cup X_{cluster} \cup X_{network}$ .

Cassandra, a distributed key-value database, is the application running on top of the network slicing used in this study. Using the *cassandra-stress* [34] tool, we generated logs of the reading and writing operations in this database to build a set of metrics  $X_{application}$ . *cassandra-stress* presents different performance metrics for reading and writing operations in the database, such as latency, operation rate, errors, and others.

The  $X_{cluster}$  represents the statistics collected from the underlying infrastructure that supports the network slice execution. The metrics collected refer to the consumption of the Central Processing Unit (CPU), Random-Access Memory (RAM), and Input/Output (I/O) operations required by the network slice from the infrastructure. We also used the NetData monitoring framework to collect CPU, RAM, and I/O metrics, as well as other metrics related to computational resources.

The  $X_{network}$  variable refers to the generic statistics collected from the network interfaces of each distributed entity comprising the network slice. We employ a native network metrics extractor (NetData) to gather statistics on the volume of transmitted, received, and lost data across the various interfaces that support the Cassandra service. These individual network interface statistics are then aggregated based on their respective timestamps.

The  $Y$  service-level metrics for the Cassandra application are defined by the mean latency of the Read (R) and Write (W) operations. This latency represents the time taken by the Cassandra application to complete R and W operations when induced by *cassandra-stress*. During the training phase of DNNs and basic ML algorithms, we extract features such as response time, errors, operations per second, and other relevant metrics from  $X_{application}$  to construct the dataset. Conversely, in the  $Y$  testing phase, our objective is to estimate the mean latency of various operations in Cassandra by considering only the generic infrastructure metrics of the testbeds.

Step ② involves training using various algorithms based on ML for time-series regression. We employ basic ML and DNNs to leverage their potential in handling complex datasets that lack linear relaxations, exhibit high-dimensional characteristics, and seamlessly adapt to new scenarios, thereby facilitating knowledge transfer.

Step ③ represents the search and adjustment of the hyperparameters for each DNN. At this stage, the hyperopt [35] tool uses Bayesian methods to find better parameters for training DNNs, such as learning rate, batch size, and epochs. At the end of this phase, the models are trained with the best hyperparameters and exported to enable inference through the SFI2 AI Management functional block.

The ④ step involves the empirical comparison of basic ML algorithms and DNNs using appropriate metrics for regression problems, such as MAE, MSE, and MAPE. MAE represents the average of the absolute differences between the predicted and actual values, while MSE is the mean of the squares of these differences, thereby emphasizing larger deviations in evaluating regression models. MAPE is a measure of relative error that expresses the difference between actual values and those predicted by a regression model as a percentage. MAPE is independent of the data scale, making it suitable for comparing the accuracy of regression models.

In step ⑤, we train the models and fine-tune the hyperparameters before integrating the trained models into SFI2 AI Management. The SFI2 architecture receives these models and makes them available for future SLA forecasting using the Predictor API.

In step ⑥, the Application Programming Interface (API) of the SFI2 AI Management block can handle some instances of  $X_{application,cluster,network}$  and return a possible condition of the QoE  $Y$  in which the network slice is conditioned online, allowing us to evaluate whether the SLA is being honored.



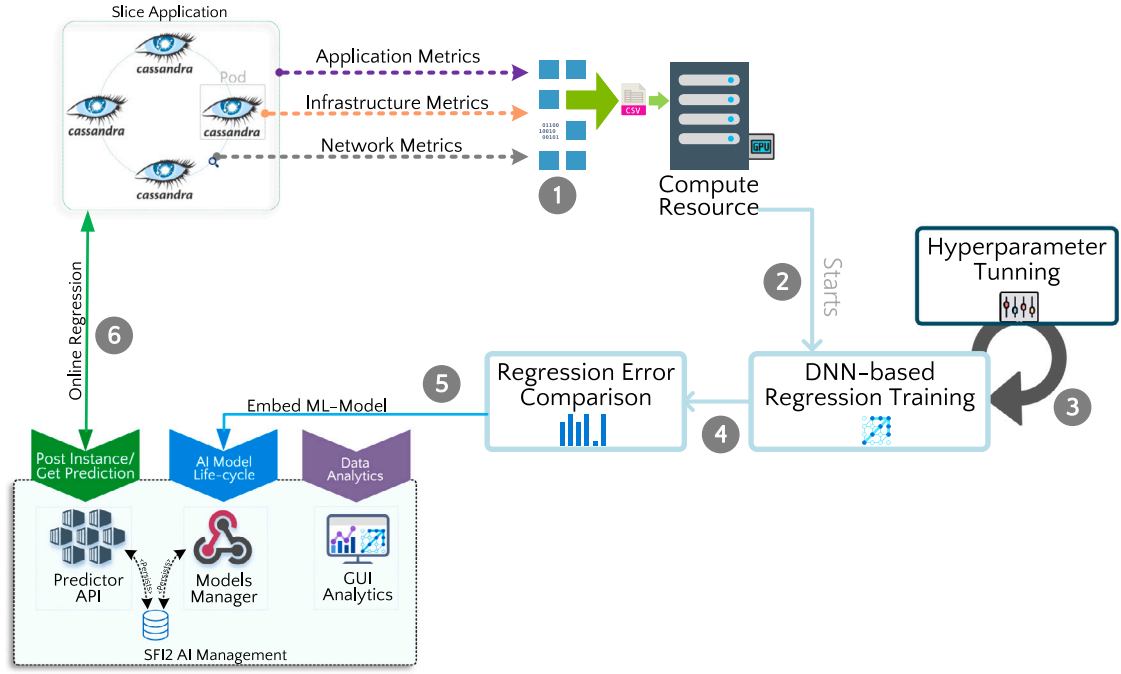


Fig. 2. Proposed method: Dataset generation combined with DNNs training and test flow.

### 3.3. Service infrastructure statistics and service-level metrics

In this section, we detail the set of input features  $X = X_{application} \cup X_{cluster} \cup X_{network}$  and denote the response variable  $Y$ . The  $X_{application}$  statistics refer to the volume of operations performed on the database at a given timestamp, errors, and lines written per second for the Cassandra application. We extracted these metrics from the `cassandra-stress` utility, thus linking each record of these statistics to the corresponding timestamp. The  $X_{cluster}$  statistics are metrics related to the consumption of the computational resources of each node that hosts the Cassandra application containers. These statistics are CPU consumption, RAM memory, and interrupts of the host machines collected through NetData.

The  $X_{network}$  metric refers to the consumption of network resources by each computing node and container in the Cassandra application. Specifically, network statistics refer to the volume of data sent, errors received from the containers that run the Cassandra ring components, and computing nodes that host the application. These statistics were collected every second using NetData.

The  $Y$  metric refers to the QoE experienced by a user when operating on a distributed database. This metric considers the average latency of write (W) and read (R) operations in Cassandra deployed on the testbed. The response variable metric  $Y$  was measured using the `Cassandra-stress` utility, and we linked each latency per second to a given timestamp. All write and read operations of our dataset generation framework in distributed testbeds did not take into account application caches in any of the replicas.

We model the collection of these metrics as  $X$  and  $Y$  time series so that our objective is to estimate the average of the W and R operations in the Cassandra application deployed by the SF12 Architecture over different testbeds using a regression problem in supervised learning [36]. Then, we estimate  $Y_t$  using machine learning algorithms that learn from the statistics  $X_t = [X_{1t}, \dots, X_{dt}]$ . Thus, we find a model  $M : X_t \rightarrow \hat{Y}_t$ , where  $\hat{Y}_t$  optimally approximates  $Y_t$  for a given  $X_t$ .

We used different ML models and DNNs model architectures to solve the regression problem. We used DT, RF, KNN, LSTM, CatBoost, and XGBoost. Among them we use DNNs such as: FCN [37], FCNPlus [38], ResNet [38], ResNetPlus [38], ResCNN [38], TCN [39], Inception-Time [40], InceptionTimePlus [40], OmniScaleCNN [41], XCM [38],

and XCMPlus [38] that are implemented in the framework Fastai and tsai [38,42].

Table 2 provides an overview of the employed DNNs structure. The Trainable Parameters column indicates the number of parameters that can be adjusted during training, including the weights and biases. The total number of Layers refers to the overall number of layers in the model, encompassing convolutional, pooling, and fully connected layers. The kernel Sizes describe the dimensions of the filters used in the convolutional layers. Pooling highlights the presence of pooling layers, which help reduce data dimensionality. Mult-Adds (M) represents the computational cost measured in millions of multiplication and addition operations. Finally, the Estimated Size (MB) estimates the model's size in megabytes, reflecting the number of parameters and required storage.

The choice of these neural networks aimed to fulfill our objective of empirically comparing the performance of DNNs in estimating the SLA compliance. We also used the Optuna [43] framework to optimize the hyperparameters of the DNNs for comparison and employed the Tree-Structured Parzen Estimator (TPE)-based algorithm [44]. Our optimizer sought to find the optimal parameters according to the search space, as shown in Table 3.

### 3.4. Dataset generation

To understand how ML algorithms perform in real testbeds, we propose a dataset that generates workloads using a periodic-load pattern and a collection framework. Using `cassandra-stress`, we generated Write and Read requests for the application deployed on the Future Internet Brazilian Environment for Experimentation New Generation (FIBRE-NG) and Fabric testbeds. These requests follow a Poisson process where the request rate adheres to a sinusoidal function, starting with a level  $P_s$  and amplitude  $P_A$  until 500k lines are written or read from the Cassandra application. The initial Cassandra parameters were defined as follows: consistency level set to *quorum*, replica factor of 2, and 256 tokens.


In Figs. 3, and 4, the workload pattern represents the number of `cassandra-stress` processes created according to a sinusoidal function over simulation time. The traces in Figs. 3 and 4 refer to the instances of `cassandra-stress`, generating requests of different

**Table 2**  
Detailed structure of employed DNN models.


Model	Trainable parameters	Total layers	Kernel sizes	Mult-Adds (M)	Estimated size (MB)
FCN	285,446	17	(7), (5), (3)	14.47	2.18
FCNPlus	285,446	18	(7), (5), (3)	14.47	2.18
ResNet	490,758	53	(7), (5), (3), (1) × 3	24.86	3.74
ResNetPlus	490,758	56	(7), (5), (3), (1) × 3	24.86	3.74
ResCNN	268,551	33	(7), (5), (3), (1), (3) × 3	13.58	2.05
TCN	71,406	95	(7) × 2, (1), (7) × 13	1.88	0.54
InceptionTime	460,038	90	(1), (39), (19), (9) × 6	23.32	3.51
InceptionTimePlus	460,038	124	(1), (39), (19), (9) × 6	23.32	3.51
OmniScaleCNN	5,239,596	68	(1), (2), (3), (5), (7), (11) × 3, (1), (2)	266.26	39.97
XCM	328,584	29	(51), (1), (51)	24.64	2.51
XCMPlus	328,584	30	(51), (1), (51)	24.64	2.51

**Table 3**  
Optuna search space for each regression model (DNNs).

Hyperparameter	Search space
Batch size	8, 16, 32
Learning Rate (LR)	0.1, 0.01, 0.001
Epochs	20, 50, 100
Patience	5, 10, 50
Optimizer	Adam, SGD
# of Layers	1, 2, 3, 4, 5
Hidden size	50, 100, 200
Bidirectional	True, False




**Fig. 3.** FIBRE-NG Testbed — Generating the traces.




**Fig. 4.** Fabric Testbed — Generating the traces.

types, such as **Write** and **Read**, to create application metrics. Whenever a new `cassandra-stress` process is created, it triggers requests (W or R) to the Cassandra application running on the testbeds. We empirically define the parameters of the sinusoidal function as  $f(t) = 22.5 + \frac{45}{2} \sin\left(\frac{2\pi t}{T}\right)$ , where  $t$  represents time and  $T$  is the period of the function.

This function models a wave that oscillates between 0 and 45, with an amplitude ( $P_A$ ) of  $\frac{45}{2}$  units and an average value ( $P_s$ ) of 22.5, providing an adequate representation for the desired variability in the generated processes. These values were defined empirically because of the restriction of computational resources for `cassandra-stress`



**Fig. 5.** FIBRE-NG: Write operation.



**Fig. 6.** FIBRE-NG: Read operation.

(container), and high values of  $P_s$  and  $P_A$  imply high resource consumption and can lead to container failure, thereby damaging the creation of the dataset. Both Write and Read operations were performed in Cassandra after the warm-up process.

Despite the Poisson model has limitations, particularly in representing peak load conditions, as it focuses on typical traffic rather than maximum loads. The dataset was generated in a live production environment using a synthetic workload application. The models were trained with realistic live production data, replicating actual network behaviors. This environment, with interconnected distributed nodes on a large-scale testbed, could experiences unexpected traffic patterns and anomalies.

This method involves forecasting the next data point based on historical observations up to the current time and making continuous predictions as new data become available. The set of data generated from the Write and Read operations, as well as the training and testing splits, is shown in Figs. 5, 6, 7, and 8, where the operation latency is ms on the y-axis, and the timestamp of the experiment is on the x-axis.

To adapt our structured numerical dataset for DNNs, specifically for Convolutional Neural Networks (CNNs), we applied a sliding window transformation using the Sliding Window method from the `tsai` library. Given a dataset with  $n$  variables (merged monitored metrics),




Fig. 7. Fabric: Write operation.




Fig. 8. Fabric: Read operation.

where the last column represents the target variable (Cassandra latency), we constructed overlapping windows of length  $w = 50$  with a stride of  $s = 1$ . Formally, let  $\mathbf{X} \in \mathbb{R}^{m \times w \times (n-1)}$  be the input tensor and  $\mathbf{y} \in \mathbb{R}^{m \times 1}$  be the corresponding target values, where  $m$  is the number of generated windows. Each window  $\mathbf{X}_i$  is defined as:

$$\mathbf{X}_i = [\mathbf{x}_i, \mathbf{x}_{i+1}, \dots, \mathbf{x}_{i+w-1}], \quad \mathbf{y}_i = \mathbf{x}_{i+w}^{\text{target}} \quad (1)$$

where  $\mathbf{x}_j \in \mathbb{R}^{n-1}$  represents the feature values at time step  $j$ . This transformation allows CNNs to extract spatial and temporal patterns effectively, treating each window as a structured input similar to an image. The sliding mechanism ensures that local dependencies are captured while enabling the model to generalize across different time segments. To evaluate model performance, we employed a time-based split using the Time Splitter function, reserving a predefined number of samples for testing.

#### 4. Experiments and model computation

Our experiments sought to validate the behavior of DNNs in estimating SLA compliance in a nationwide network slice deployed through the SFI2 Orchestrator. Initially, we deployed the application to the testbeds and started the workload tests to generate datasets for different metrics. For each testbed, we seek to identify the Cassandra application's behavior and SLA and understand whether DNNs or basic ML algorithms can generalize predictions across geographically distributed testbeds that experience production-ready network conditions.

##### 4.1. Testbeds

Fig. 9 shows the experimental setup used for our evaluation. What stands out is the Cassandra ring deployed on different computing nodes is spread across each testbed FIBRE-NG [14] and Fabric [15]. We collected and processed the monitoring metrics, which are the features  $X = X_{\text{application}} \cup X_{\text{cluster}} \cup X_{\text{network}}$ , feeding the Predictor API that applies training with the different regression models.

To collect the application metrics, we employed a sensor and load-generating node. The load-generating pipeline shown in Fig. 9 executed

read and write requests against the Cassandra ring (step ①), whereas the sensor node generated flat requests simultaneously without altering the workload pattern (step ②). Thus, we successfully gathered the necessary metrics of the application from the server side, which were subsequently compiled (step ③) into a dataset. Finally, the dataset is uploaded to the SLA compliance training and validation framework (step ④).

We deployed our network slice with the Cassandra application on two nationwide testbeds, FIBRE-NG and Fabric, considering the computational nodes in different geographic and intercontinental locations, such as the European Organization for Nuclear Research (CERN) node, presented in Table 4, which shows the rest of each testbed where we deployed the Cassandra application. This experimental setup aims to validate how DNNs deal with traces generated at nodes that transmit data in a production-ready network.

##### 4.2. Model training

In the experimental training and validation phase of the basic ML algorithms and DNNs, we used RTX 4060ti 16 Gb GPU hardware with an CPU Intel(R) Core(TM) i5-4430 CPU @3.00 GHz with 32 GB of RAM. Furthermore, we used the PyTorch framework together with FastAi [42] and tsai [38] tools to build the DNNs training setup. All models were trained and validated with no seed locks and with tuned hyperparameters, and we collected ten (10) samples of training time and performance metrics. The choice of DNNs and basic ML models aims to represent different model architectures for generalization and in-depth analysis.

Our training and test pipeline involves data ingestion, timestamping, and indexing, followed by splitting into training and testing subsets. Data transformation includes resampling, interpolation, and normalization, with feature selection and separate scaling of the target variable. The Sliding Window method generates input-output pairs for time series forecasting, while TSDatasets and TSDataLoaders apply further transformations. Hyperopt optimizes hyperparameters like batch size and learning rate. The model is trained using Learner with early stopping to prevent overfitting, and performance metrics are recorded for evaluation, ensuring robust data preparation and optimization.

#### 5. Evaluating results

To validate our contribution, we initially analyze the impact of different network factors on our case study network slicing application response time. We also discuss aspects of Model Tuning for our employed neural network. Later, we evaluate different ML approaches on a production-ready network and assess the feasibility of our dataset generation method, which expresses real network conditions and enables fitting and training models to forecast SLA violations on network slicing architectures.

##### 5.1. Impact analysis in a large-scale network

To assess the impact of realistic and production-ready network metrics on our experimental deployment, we used a chaos engineering tool to simulate latency with jitter and packet loss [45]. We systematically applied jitter ranging from 1 ms to 10 ms, following a uniform distribution, and introduced a packet loss between 1% and 10% in our experimental sliced application. For token management, we implemented two distinct slices containing our distributed database application by varying the number of tokens from 32 to 256.

We employed three-way Analysis of Variance (ANOVA) to evaluate the effects of three independent factors and their interactions on a dependent variable. The factors are: (1) Fixed Latency with Jitter induction with levels 1 ms and 10 ms, (2) Network Packet Loss with levels 1% and 10%, and (3) Cassandra Tokens with levels 32 and 256. The response variable analyzed is the Cassandra Write and Read

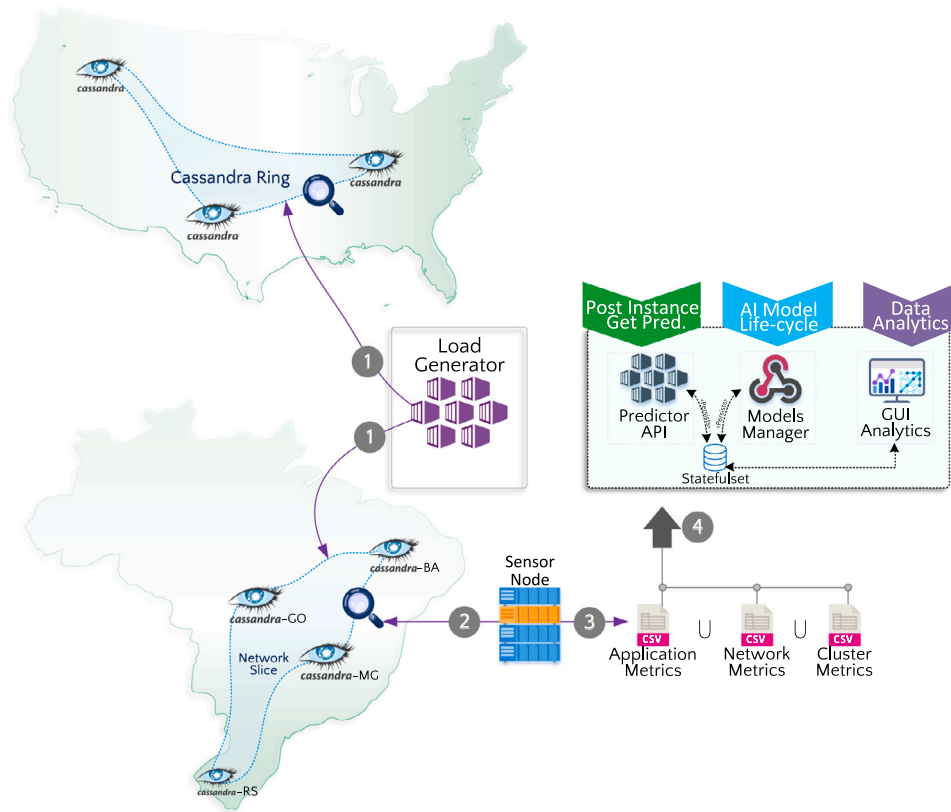


Fig. 9. Deployment of Cassandra application on nationwide testbeds.

**Table 4**  
Testbeds compute nodes hosting Cassandra services (containers) in a distributed manner.

Testbed	Pod name	Node name	Location
FIBRE-NG	cassandra-0	WHX-SC	Santa Catarina
	cassandra-1	WHX-RS	Rio Grande do Sul
	cassandra-2	WHX-PB	Paraíba
	loadgen	WHX-RN	Rio Grande do Norte
Fabric	cassandra-0	Great Plains Network (GPN)	Kansas City, MO
	cassandra-1	CERN	France
	cassandra-2	The University of Utah	Salt Lake City, UT
	loadgen	Rutgers University	Jersey City, NJ

**Table 5**  
Experimental factors influence on write operations.

Source of variation	F value	P value
Network Delay	7.58245	0.00612
Network Loss	580.45267	<0.0001
Cassandra Tokens	0.25329	0.615
Network Delay × Network Loss	4.32445	0.03811
Network Delay × Cassandra Tokens	0.45758	0.49909
Network Loss × Cassandra Tokens	1.24777	0.26455
Network Delay × Network Loss × Cassandra Tokens	1.77435	0.18349

Latency. This method tests the main effects of each factor, the two-way interactions between factors, and the three-way interaction, using the F-statistic and corresponding  $p$ -value. A  $p$ -value of  $p < 0.05$  indicates that the effect or interaction is statistically significant.

The three-way ANOVA results as Table 5 show that for Write operations, Network Delay (Jitter) ( $F(1,472) = 7.58$ ,  $p = 0.00612$ ) and Network Packet Loss ( $F(1,472) = 580.45$ ,  $p < 0.0001$ ) have significant effects on latency, while Cassandra Tokens ( $F(1,472) = 0.25$ ,  $p = 0.615$ ) do not. The interaction between Network Delay and Network Packet Loss is also significant ( $F(1,472) = 4.32$ ,  $p = 0.03811$ ), suggesting non-additive effects. However, interactions involving Cassandra Tokens and

the three-way interaction are not significant ( $p > 0.05$ ). These results emphasize that network-related factors are the primary contributors to latency during Write operations.

The results in Fig. 10 that high network latency (10 ms) combined with high packet loss (0.1) significantly degrades performance, as seen in the red line dropping from 9000 to 8000. In contrast, low latency (1 ms) with high packet loss starts at around 7000, with a slight increase. Configurations with low packet loss (0.01) maintain stable performance despite increased latency. ANOVA confirms these findings, with network delay and loss showing high statistical significance ( $P < 0.0001$ ). These results highlight the importance of optimizing network conditions for enhancing the efficiency and reliability of distributed database systems like Cassandra.

The three-way ANOVA results for Read operations as Table 6 indicate that both Network Delay ( $F(1,473) = 4.19$ ,  $p = 0.04124$ ) and Network Packet Loss ( $F(1,473) = 172.78$ ,  $p < 0.0001$ ) significantly affect latency. Additionally, Cassandra Tokens shows a marginally significant effect ( $F(1,473) = 3.81$ ,  $p = 0.0515$ ), suggesting a potential influence on latency that may be worth further exploration. The interaction between Network Delay and Network Packet Loss is also nearly significant ( $F(1,473) = 3.58$ ,  $p = 0.05896$ ), while the interaction between Network Delay and Cassandra Tokens is not significant ( $F(1,473) =$






Fig. 10. Impact of Network Latency, Cassandra Tokens and Packet Loss on Cassandra Write Operations.

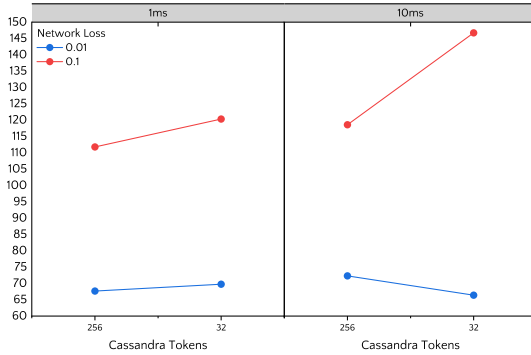


Fig. 11. Impact of Network Latency, Cassandra Tokens and Packet Loss on Cassandra Read Operations.

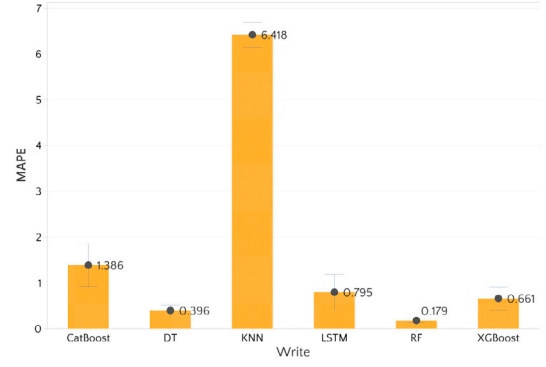
Table 6

Experimental factors influence on read operations.

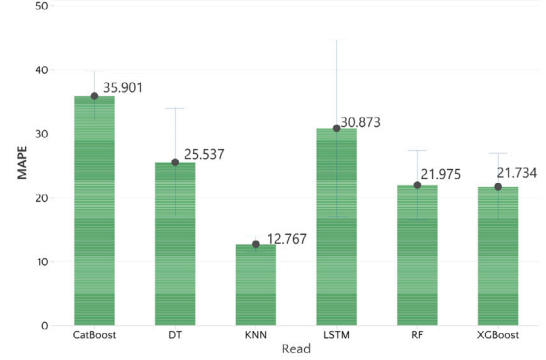
Source of variation	F value	p value
Network Delay	4.19	0.04124
Network Packet Loss	172.78	<0.0001
Cassandra Tokens	3.81	0.0515
Network Delay × Network Packet Loss	3.58	0.05896
Network Delay × Cassandra Tokens	0.48	0.491
Network Loss × Cassandra Tokens	5.79	0.01648
Network Delay × Network Loss × Cassandra Tokens	2.69	0.10178

0.48,  $p = 0.491$ ). Furthermore, the interaction between Network Loss and Cassandra Tokens is significant ( $F(1, 473) = 5.79$ ,  $p = 0.01648$ ), suggesting that these two factors jointly influence latency. The three-way interaction between Network Delay, Network Loss, and Cassandra Tokens is not significant ( $F(1, 473) = 2.69$ ,  $p = 0.10178$ ). These findings highlight the dominant influence of network-related factors on latency during Read operations, with Cassandra Tokens potentially having a marginal effect and interactions between network factors playing a role.

In Fig. 11 suggests that high network latency (10 ms) combined with high packet loss (0.1) significantly degrades performance, with a notable decline observed in the graph. Conversely, lower latency (1 ms) with high packet loss exhibits a slight performance increase. Moreover, configurations with low packet loss (0.01) demonstrate minimal impact from increased latency, indicating stability in performance. These trends are corroborated by ANOVA results, highlighting the statistical significance of network delay ( $P = 0.00612$ ) and network loss ( $P < 0.0001$ ), as well as their interaction ( $P = 0.03811$ ). These findings underscore the importance of optimizing both latency and packet loss to enhance the efficiency and reliability of distributed database systems like Cassandra.



(a) FIBRE-NG: MAPE (Write).



(b) FIBRE-NG: MAPE (Read).

Fig. 12. MAPE results for Write and Read operations on FIBRE-NG.

## 5.2. Model tuning

In this section, we present the results of the hyperparameter optimization of DNNs and the performance in estimating SLA for each model for the different operations (W and R) and testbeds (FIBRE-NG and Fabric). Therefore, we conducted hyperparameter optimization for the four constructed datasets ‘fibre-read.csv’, ‘fibre-write.csv’, ‘fabric-read.csv’, and ‘fabric-write.csv’.

We summarize the hyperparameters found by Optuna for the FIBRE-NG testbed in Tables 7, and 8. Furthermore, we tuned the hyperparameters of the same models for the dataset generated from the Fabric testbed, as shown in Tables 9, 10. With these adjusted values, we proceeded with an empirical evaluation of the performance of these models for predicting the SLA in each testbed.

## 5.3. Basic model performance

Through our dataset construction framework on sliced testbeds, we aggregate the metrics  $X = X_{application} \cup X_{cluster} \cup X_{network} \rightarrow Y$  and configure the datasets generated to submit it to the training and assessment process. Thus, we empirically use the division 80% for training and 20% for testing and build the regression model aiming at the one-step-ahead prediction method that, given a statistic  $i$  of  $X_i$ , the model can estimate the operation latency (W or R) in the next step.

As Fig. 12-a, the Write operation on the FIBRE-NG testbed, and RF presented the lowest mean MAPE (0.17), indicating the best predictive accuracy, followed by DT (0.39). The KNN had the worst performance, with an average MAPE of 6.41, which is significantly higher than that of the others. Models such as CatBoost, LSTM, and XGBoost have intermediate performance but are still superior to KNN. In addition, RF demonstrated greater stability, with the lowest standard deviation (0.064).

**Table 7**

Write Dataset: Hyperparameter tuning values for the FIBRE-NG testbed.

Model	Batch size	Learning rate	Epochs	Patience	Optimizer	Layers	Hidden size	Bidirectional
FCN	8	0.1	20	10	Adam	2	100	False
FCNPlus	16	0.1	20	50	Adam	4	200	False
InceptionTime	16	0.1	100	50	Adam	4	100	False
InceptionTimePlus	16	0.1	100	50	Adam	1	50	True
OmniScaleCNN	32	0.1	20	50	Adam	2	100	True
ResCNN	32	0.1	20	10	Adam	5	100	False
ResNet	32	0.1	50	50	Adam	2	100	False
ResNetPlus	32	0.1	20	5	Adam	5	50	False
TCN	32	0.01	100	50	Adam	1	200	False
XCM	16	0.01	50	50	Adam	1	200	False
XCMPlus	8	0.1	100	50	Adam	3	100	True

**Table 8**

Read Dataset: Hyperparameter tuning values for the FIBRE-NG testbed.

Model	Batch size	Learning rate	Epochs	Patience	Optimizer	Layers	Hidden size	Bidirectional
FCN	16	0.1	20	10	Adam	3	200	True
FCNPlus	16	0.1	20	50	Adam	4	50	False
InceptionTime	8	0.01	20	10	Adam	5	100	True
InceptionTimePlus	16	0.1	20	50	Adam	2	200	False
OmniScaleCNN	8	0.1	50	50	Adam	2	50	True
ResCNN	16	0.1	20	10	Adam	3	100	False
ResNet	16	0.1	20	10	Adam	1	200	True
ResNetPlus	8	0.1	20	50	Adam	2	50	False
TCN	32	0.01	50	5	Adam	4	50	False
XCM	8	0.01	20	10	Adam	3	200	False
XCMPlus	16	0.01	50	50	Adam	3	50	True

**Table 9**

Write Dataset: Hyperparameter tuning values for the Fabric Testbed.

Model	Batch size	Learning Rate (LR)	Epochs	Patience	Optimizer	# of Layers	Hidden size	Bidirectional
FCN	32	0.01	20	10	Adam	3	50	True
FCNPlus	32	0.1	100	50	Adam	5	100	True
InceptionTime	32	0.1	50	50	Adam	2	200	False
InceptionTimePlus	16	0.1	100	50	Adam	4	50	True
OmniScaleCNN	16	0.01	50	10	Adam	4	100	True
ResCNN	32	0.01	50	50	Adam	5	100	False
ResNet	32	0.01	50	50	Adam	4	50	True
ResNetPlus	16	0.1	20	10	Adam	3	200	True
TCN	8	0.001	100	50	Adam	2	50	False
XCM	16	0.1	50	50	Adam	3	50	False
XCMPlus	16	0.1	100	50	Adam	1	50	True

**Table 10**

Read Dataset: Hyperparameter tuning values for the fabric testbed.

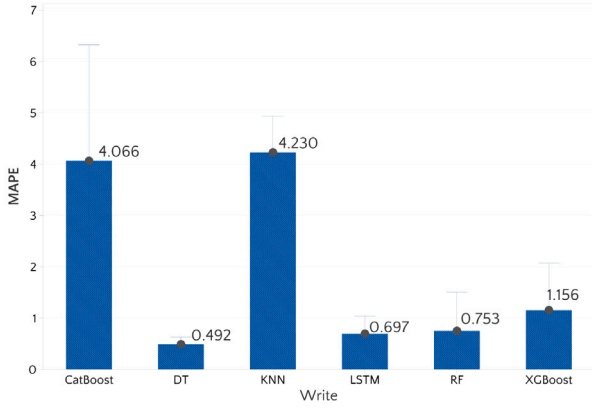
Model	Batch size	LR	Epochs	Patience	Optimizer	# Layers	Hidden size	Bidirectional
FCN	32	0.1	20	10	Adam	3	100	False
FCNPlus	8	0.1	20	10	Adam	3	200	True
InceptionTime	32	0.1	100	10	Adam	5	50	True
InceptionTimePlus	16	0.1	100	50	Adam	2	50	False
OmniScaleCNN	8	0.001	20	50	SGD	4	100	False
ResCNN	32	0.1	20	10	SGD	2	100	True
ResNet	32	0.1	20	50	Adam	4	100	True
ResNetPlus	8	0.1	20	5	Adam	4	200	False
TCN	32	0.01	20	10	Adam	2	100	True
XCM	8	0.1	100	10	SGD	3	50	True
XCMPlus	8	0.1	100	10	SGD	5	200	True

In the Read operation, as shown in Fig. 12-b on the FIBRE-NG testbed, KNN obtained the lowest mean MAPE (12.77), indicating the best predictive accuracy. In contrast, CatBoost (35.90) and LSTM (30.87) presented the highest errors with high variability, as evidenced by the high standard deviation of LSTM (19.34). RF and XGBoost had intermediate performances, with mean MAPEs of 21.98 and 21.73, respectively. Meanwhile, DT showed a mean error of 25.54, with high dispersion.


During the writing operation, as shown in Fig. 13-a, on the testbed Fabric, the DT presented the lowest mean MAPE (0.49), indicating

the best predictive accuracy, followed by LSTM (0.69) and RF (0.75). XGBoost had a slightly higher error (1.16) but was still lower than CatBoost (4.07) and KNN (4.23), which had the worst performance. In addition, DT showed the lowest standard deviation (0.19), suggesting greater stability.

In the Read operation, as Fig. 13-b, on the testbed Fabric, RF presented the lowest mean MAPE (4.06), closely followed by XGBoost (4.09) and CatBoost (4.03), indicating very similar performances. DT had a higher error (5.64), while KNN obtained the worst result (9.29), with the highest standard deviation (0.22), demonstrating



(a) Fabric: MAPE (Write).



(b) Fabric: MAPE (Read).

Fig. 13. MAPE results for Write and Read operations on Fabric.

low precision. LSTM showed an intermediate MAPE (4.85), but with greater variability (0.53). Thus, RF, XGBoost, and CatBoost stand out as the most effective approaches for predicting the latency in the Read operation of the distributed Cassandra on the testbed Fabric.

These experiments with basic ML algorithms for forecasting highlight the contribution of our study, which discusses the suitability of laboratory-trained algorithms in real-world scenarios and production networks with geographically distributed nodes. In the four analyzed scenarios, RF and DT demonstrated the lowest mean MAPE, indicating higher predictive accuracy. In the Write operation on the Fabric testbed, DT had the best performance (0.49), while in the Read operation on the same testbed, RF stood out (4.06). On FIBRE-NG, RF and KNN were more accurate in the Read operation (21.98 and 12.77, respectively), whereas DT had the lowest error in the Write operation (0.40). In contrast, KNN and CatBoost exhibited the worst performances in various scenarios. These results suggest that tree-based models, especially RF and DT, are more effective for forecasting in sliced testbeds.

#### 5.4. DNN-based model performance

We computed the DNNs training time demands for Write and Read operations on the FIBRE-NG and Fabric testbeds and compared the results in Fig. 14. Therefore, it is possible to infer that with a confidence level of 95%, the read datasets generated in different testbeds require the same training time. In contrast, we noticed a difference in the training time between the Read and Write operations owing to the size of the generated dataset and the variations in network conditions experienced by each test workload.

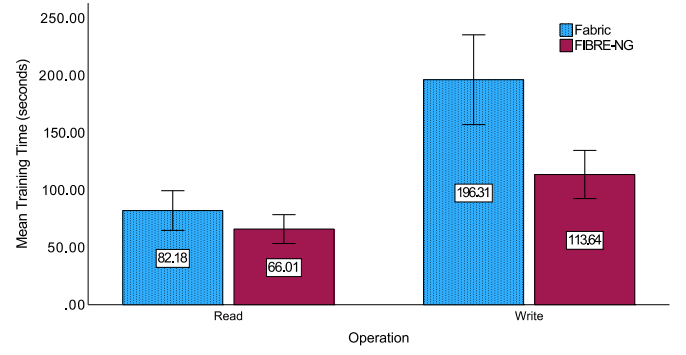


Fig. 14. Training Time for different Operations (W and R) on Testbeds.

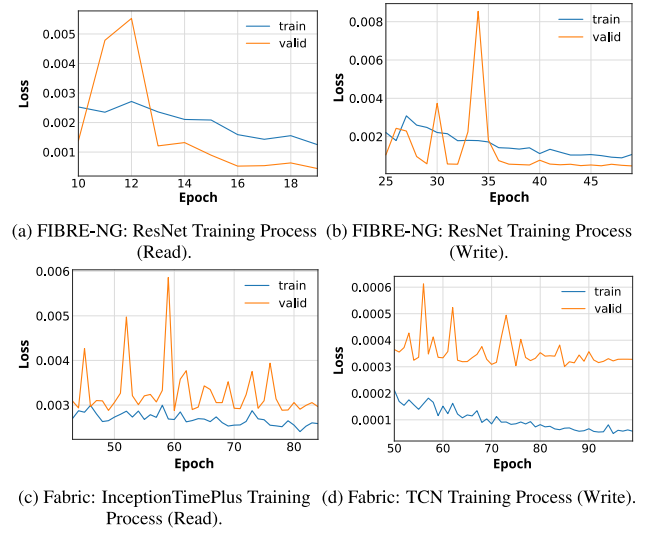


Fig. 15. Training and Test behavior for Read and Write datasets on FIBRE-NG and Fabric Testbeds.

We deepened our analysis by observing the generalization capacity of the DNNs used to predict the Write and Read operations latency in Cassandra deployed on the FIBRE-NG and Fabric Testbeds. We present in Fig. 14 the training and validation behaviors of only the best DNNs for the two operations (W and R) and testbeds considering the MAE metric and the measured values in Tables 11 and 12.

We choose to employ the MAE metric rather than RMSE or MSE for regression on the Read and Write latency owing to its robustness to outlier handling and ease of interpretation. Additionally, MAE shares the same unit of measurement as the dependent variable, facilitating a more intuitive comprehension of its meaning and magnitude.

For the read and write datasets from the FIBRE-NG testbed shown in Fig. 15-a and -b, we observed the training behavior of ResNet DNN. From Fig. 15-c, which refers to DNN InceptionTimePlus for the read dataset in the Fabric testbed, there was a subtle drop in the training and test losses, indicating that DNNs could extract patterns from the time series for prediction. As with Fig. 15-d, which refers to the dataset written in the Fabric testbed, there is a visual indication that the loss decreases as the epochs advance.

The results presented in Tables 13 and 14 summarize the performance of DNNs in the context of FIBRE-NG, utilizing the RMSE metric. The InceptionTimePlus network stands out with the lowest RMSE for both reading (0.011) and writing (0.016), indicating superior performance compared to other networks. The ResNet also shows competitive results, with RMSE values of 0.010 for reading and 0.012 for writing. In contrast, the XCM and OmniScaleCNN exhibit the highest RMSE

**Table 11**  
Summarizing the DNNs performance for FABRIC using MAE.

Neural network	Read mean	Read StDev	Write mean	Write StDev
FCN	0.040	0.000	0.016	0.005
FCNPlus	0.042	0.004	0.012	0.004
InceptionTime	0.041	0.003	0.020	0.008
InceptionTimePlus	0.040	0.000	0.030	0.017
OmniScaleCNN	0.046	0.005	0.027	0.018
ResCNN	0.043	0.005	0.021	0.007
ResNet	0.040	0.000	0.059	0.033
ResNetPlus	0.048	0.006	0.024	0.007
TCN	0.040	0.000	0.012	0.004
XCM	0.054	0.014	0.027	0.018

**Table 12**  
Summarizing the DNNs performance for FIBRE-NG using MAE.

Neural network	Read mean	Read StDev	Write mean	Write StDev
FCN	0.017	0.008	0.021	0.007
FCNPlus	0.013	0.005	0.014	0.005
InceptionTime	0.013	0.005	0.024	0.027
InceptionTimePlus	0.011	0.003	0.016	0.007
OmniScaleCNN	0.020	0.009	0.035	0.025
ResCNN	0.021	0.019	0.026	0.013
ResNet	0.010	0.000	0.012	0.004
ResNetPlus	0.014	0.005	0.017	0.005
TCN	0.019	0.007	0.015	0.005
XCM	0.025	0.007	0.035	0.012

**Table 13**  
Summarizing the DNNs performance for FABRIC using RMSE.

Neural network	Read mean	Read StDev	Write mean	Write StDev
FCN	0.050	0.000	0.020	0.000
FCNPlus	0.052	0.004	0.019	0.003
InceptionTime	0.055	0.005	0.025	0.010
InceptionTimePlus	0.053	0.005	0.034	0.016
OmniScaleCNN	0.059	0.003	0.033	0.021
ResCNN	0.054	0.005	0.024	0.007
ResNet	0.051	0.003	0.066	0.038
ResNetPlus	0.061	0.010	0.030	0.012
TCN	0.050	0.000	0.021	0.003
XCM	0.065	0.014	0.020	0.005
XCMPlus	0.065	0.009	0.026	0.005

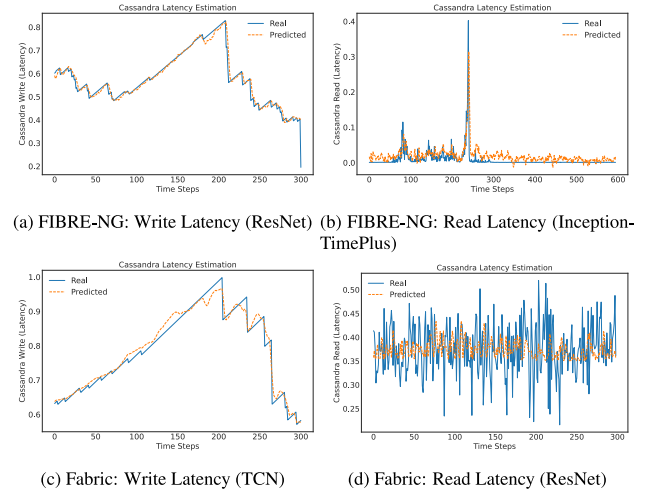
**Table 14**  
Summarizing the DNNs performance for FIBRE-NG using RMSE.

Neural network	Read mean	Read StDev	Write mean	Write StDev
FCN	0.017	0.008	0.021	0.007
FCNPlus	0.013	0.005	0.014	0.005
InceptionTime	0.013	0.005	0.024	0.027
InceptionTimePlus	0.011	0.003	0.016	0.007
OmniScaleCNN	0.020	0.009	0.035	0.025
ResCNN	0.021	0.019	0.026	0.013
ResNet	0.010	0.000	0.012	0.004
ResNetPlus	0.014	0.005	0.017	0.005
TCN	0.019	0.007	0.015	0.005
XCM	0.025	0.007	0.035	0.012

values, suggesting lower accuracy in their predictions. Additionally, the table reveals the variability of the results, reflected in the standard deviations (StDev), which vary across the different architectures, with most maintaining relatively low deviations, especially in reading tasks.

Having the DNNs learn over the epochs operating on the datasets generated in the two testbeds, it is possible to admit that the DNNs are efficient in dealing with the seasonality of a production-ready network. Then, slicing orchestrators can couple such models into their slicing management control loop and modernize the delivery of service verticals with a guaranteed SLA.

We seek to observe the behavior of DNNs in predicting Cassandra latency by contrasting the real and predicted in the test portion of the time series. Fig. 16 shows the DNNs and their prediction process



**Fig. 16.** Latency prediction for better DNNs for Write and Read operations in testbeds.

performance, considering the lowest value of the MAE metric for both the Write or Read operations presented in Tables 11 and 12. Thus, modern slicing orchestrations can adopt a threshold for the difference between actual and predicted and assess whether the network slices it manages comply with the agreed SLA.

To further our analysis, we examined the performance of DNNs for the four datasets generated in the two testbeds using the MAPE metric. As illustrated in Fig. 17-a, for the FIBRE-NG testbed, ResNet demonstrated the highest performance on the Write dataset with a MAPE of 0.024. Conversely, for the Read dataset, DNN InceptionTimePlus exhibited superior performance, as depicted in Fig. 17-b.

Furthermore, our analysis revealed that the DNN exhibiting optimal performance for the Write dataset in the Fabric testbed was FCNPlus (Fig. 18-a), achieving a MAPE of 0.015, with standard deviation serving as the criterion for resolving ties.

In the Fabric Read scenario, and according to Fig. 18-b, the MAPE showed very low and homogeneous values across the models, indicating high prediction accuracy. The low variability suggests that the prediction task for this operation was relatively simple, resulting in insignificant percentage errors. Models such as FCN, ResNet, and TCN performed practically identically, whereas XCM and XCMPlus showed slight variations, but still within a very small error range.

Considering this empirical analysis of the behavior of DNNs and the MAPE metric, it is possible to admit that DNNs are technologies that perform intelligent network slicing. It is possible to couple such capabilities into different building blocks to act at different phases of the lifecycle of a network slice. Using the MAPE metric, it is possible to have a percentage dimension of the error of DNNs that reiterates its ability to be embedded in prediction Application Programming Interfaces (APIs) based on microservices, such as SF12 Orchestration Architecture.



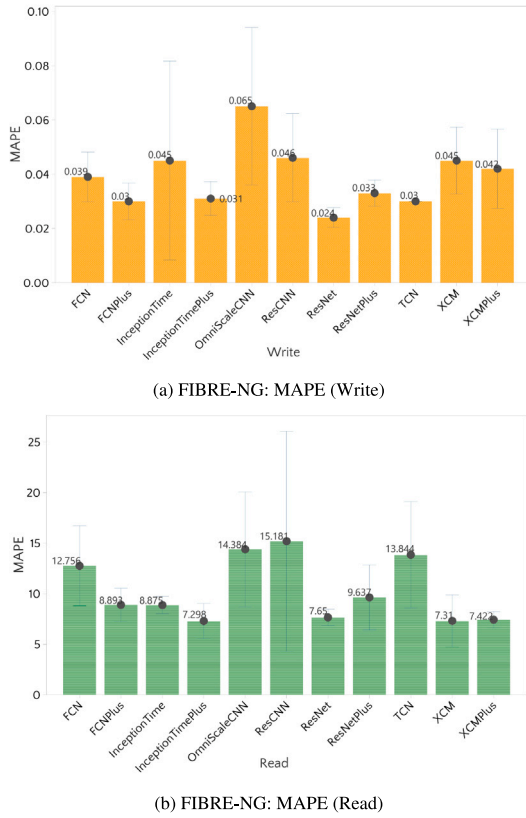


Fig. 17. Mean Absolute Percentage Error (MAPE) for FIBRE-NG in Write and Read operations.

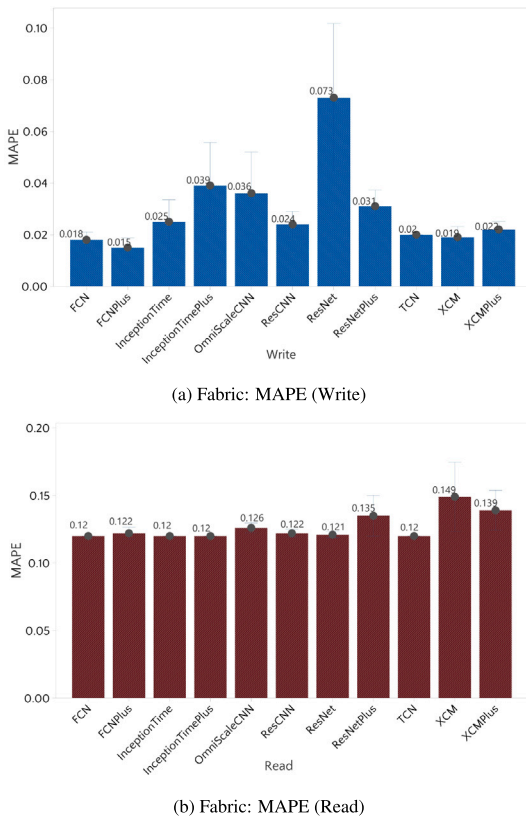


Fig. 18. Mean Absolute Percentage Error (MAPE) for Fabric in Write and Read operations.

Our approach primarily focused on evaluating a distributed database to assess the feasibility of training ML algorithms on a large-scale testbed. In addition, our method can be extended to a broader range of applications that require access to computing and network-monitoring platforms. This extension would enable the integration of diverse performance metrics, facilitating accurate application performance forecasting and supporting life-cycle decision making in network slicing architectures.

## 6. Concluding remarks

In this study, we shed light on how ML techniques, specifically DNNs and basic ML algorithms, can be jointly employed with slicing orchestration architectures to leverage and guarantee SLA for tailored applications in nationwide testbeds. To achieve this, we propose a method for generating and aggregating datasets regarding the latency of Write and Read operations in a distributed database. We found that there are approaches in the literature that combine computational intelligence for the different phases of the network slice life cycle; however, they have not yet considered how these AI techniques behave in production-ready networks deployed on nationwide testbeds.

Among the findings, we found that DNNs and even basic ML algorithms are promising technologies that can be built or natively embedded in slicing architectural building blocks to perform zero-touch orchestration in production-ready networks. Furthermore, we verified that forecasting network slicing latency with a low error rate is possible by monitoring generic and easily collected metrics related to the computing or network resources on which the network slice is deployed. Furthermore, we believe that embedding the DNN or basic ML models in SFI2 AI management to cope with stringent application vertical requirements is a promising path.

One of the constraints of this study is that it focuses on generic networks and computing metrics. We aimed to incorporate more heterogeneous metrics into the dataset construction process to assess the generalization of these metrics and achieve low error rates in our predictions. Currently, we are working on analyzing the separation of metrics to validate the impact of each on the final ability to estimate SLA conformance and the employment of Reinforcement Learning.

In addition, we plan to explore methods such as queueing theory, extreme value analysis, or bursty traffic models to better capture extreme network conditions using supervised learning methods and other DNNs and attention-based mechanisms to determine their efficacy in such contexts. Our results offer valuable insights and opportunities for the further exploration of intelligent native slicing architectures.

## CRedit authorship contribution statement

**Rodrigo Moreira:** Writing – original draft, Visualization, Validation, Software, Resources, Data curation, Conceptualization. **Rafael Pasquini:** Writing – review & editing, Validation, Formal analysis. **Joberto S.B. Martins:** Writing – review & editing, Validation, Supervision, Funding acquisition. **Tereza C. Carvalho:** Project administration, Investigation, Funding acquisition, Formal analysis. **Flávio de Oliveira Silva:** Writing – review & editing, Validation, Supervision, Project administration, Methodology, Funding acquisition.

## Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Tereza Cristina Melo de Brito Carvalho reports financial support was provided by State of Sao Paulo Research Foundation. Joberto Sergio Barbosa Martins reports financial support was provided by Anima Institute. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

We acknowledge the financial support of the Brazilian National Council for Scientific and Technological Development (CNPq), grant #421944/2021-8, and the FAPESP MCTIC/CGI, Brazil Research project 2018/23097-3 - SFI2 - Slicing Future Internet Infrastructures. The authors also thank CNPq, CAPES, FAPESP and Instituto ANIMA.

## Data availability

Data will be made available on request.

## References

- [1] Y. Liu, Y. Deng, A. Nallanathan, J. Yuan, Machine learning for 6G enhanced ultra-reliable and low-latency services, *IEEE Wirel. Commun.* 30 (2) (2023) 48–54, <http://dx.doi.org/10.1109/MWC.006.2200407>.
- [2] M. Zhang, L. Shen, X. Ma, J. Liu, Toward 6G-Enabled mobile vision analytics for immersive extended reality, *IEEE Wirel. Commun.* 30 (3) (2023) 132–138, <http://dx.doi.org/10.1109/MWC.023.2200557>.
- [3] J. Li, H. Wu, X. Huang, Q. Huang, J. Huang, X.S. Shen, Toward Reinforcement-Learning-Based intelligent network control in 6G networks, *IEEE Netw.* 37 (4) (2023) 104–111, <http://dx.doi.org/10.1109/MNET.003.2200641>.
- [4] A.T. Jawad, R. Maaloul, L. Chaari, A comprehensive survey on 6G and beyond: Enabling technologies, opportunities of machine learning and challenges, *Comput. Netw.* 237 (2023) 110085, <http://dx.doi.org/10.1016/j.comnet.2023.110085>, URL <https://www.sciencedirect.com/science/article/pii/S1389128623005303>.
- [5] R. Moreira, P.F. Rosa, R.L.A. Aguiar, F.d. Silva, NASOR: A network slicing approach for multiple autonomous systems, *Comput. Commun.* 179 (2021) 131–144, <http://dx.doi.org/10.1016/j.comcom.2021.07.028>, URL <https://www.sciencedirect.com/science/article/pii/S0140366421002917>.
- [6] S. S., S. Mishra, C. Hota, Joint QoS and energy-efficient resource allocation and scheduling in 5G network slicing, *Comput. Commun.* 202 (2023) 110–123, <http://dx.doi.org/10.1016/j.comcom.2023.02.009>, URL <https://www.sciencedirect.com/science/article/pii/S0140366423000464>.
- [7] R. Moreira, L.F. Rodrigues Moreira, F. de Oliveira Silva, An intelligent network monitoring approach for online classification of Darknet Traffic, *Comput. Electr. Eng.* 110 (2023) 108852, <http://dx.doi.org/10.1016/j.compeleceng.2023.108852>, URL <https://www.sciencedirect.com/science/article/pii/S0045790623002768>.
- [8] R. Ahmed, M.R. Mahmood, M.A. Matin, Challenges in meeting QoS requirements toward 6G wireless networks: A state of the art survey, *Trans. Emerg. Telecommun. Technol.* 34 (2) (2023) e4693, <http://dx.doi.org/10.1002/ett.4693>, arXiv: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/ett.4693> URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/ett.4693>.
- [9] M.E. Haque, F. Tariq, M.R.A. Khandaker, K.-K. Wong, Y. Zhang, A survey of scheduling in 5G URLLC and outlook for emerging 6G systems, *IEEE Access* 11 (2023) 34372–34396, <http://dx.doi.org/10.1109/ACCESS.2023.3264592>.
- [10] W. Wu, C. Zhou, M. Li, H. Wu, H. Zhou, N. Zhang, X.S. Shen, W. Zhuang, AI-native network slicing for 6G networks, *IEEE Wirel. Commun.* 29 (1) (2022) 96–103, <http://dx.doi.org/10.1109/MWC.001.2100338>.
- [11] P. Singh, Systematic review of data-centric approaches in artificial intelligence and machine learning, *Data Sci. Manag.* 6 (3) (2023) 144–157, <http://dx.doi.org/10.1016/j.dsm.2023.06.001>, URL <https://www.sciencedirect.com/science/article/pii/S2666764923000279>.
- [12] J. Gomez, E.F. Kfoury, J. Crichigno, G. Srivastava, A survey on network simulators, emulators, and testbeds used for research and education, *Comput. Netw.* 237 (2023) 110054, <http://dx.doi.org/10.1016/j.comnet.2023.110054>, URL <https://www.sciencedirect.com/science/article/pii/S1389128623004991>.
- [13] J.S.B. Martins, T.C. Carvalho, R. Moreira, C.B. Both, A. Donatti, J.H. Corrêa, J.A. Suruagy, S.L. Corrêa, A.J.G. Abelem, M.R.N. Ribeiro, J.-m.S. Nogueira, L.C.S. Magalhães, J. Wickboldt, T.C. Ferreto, R. Mello, R. Pasquini, M. Schwarz, L.N. Sampaio, D.F. Macedo, J.F. De Rezende, K.V. Cardoso, F. De Oliveira Silva, Enhancing network slicing architectures with machine learning, security, sustainability and experimental networks integration, *IEEE Access* 11 (2023) 69144–69163, <http://dx.doi.org/10.1109/ACCESS.2023.3292788>.
- [14] T. Salmato, L. Ciuffo, I. Machado, M. Salvador, M. Stanton, N. Rodriguez, A. Abelem, L. Bergesio, S. Sallent, L. Baron, Fibre-an international testbed for future internet experimentation, in: *Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos-SBRC 2014*, 2014, pp. p-969.
- [15] I. Baldin, A. Nikolich, J. Griffioen, I.I.S. Monga, K.-C. Wang, T. Lehman, P. Ruth, FABRIC: A national-scale programmable experimental network infrastructure, *IEEE Internet Comput.* 23 (6) (2019) 38–47, <http://dx.doi.org/10.1109/MIC.2019.2958545>.
- [16] H.P. Phyu, D. Naboulsi, R. Stanica, Machine learning in network slicing—A survey, *IEEE Access* 11 (2023) 39123–39153, <http://dx.doi.org/10.1109/ACCESS.2023.3267985>.
- [17] R. Moreira, T. Carvalho, F. Silva, Designing and evaluating a high-reliable and security-aware Identity and access management for slicing architectures, in: *Anais do XIV Workshop de Pesquisa Experimental da Internet do Futuro*, SBC, Porto Alegre, RS, Brasil, 2023, pp. 1–6, <http://dx.doi.org/10.5753/wpeif.2023.722>, URL <https://sol.sbc.org.br/index.php/wpeif/article/view/24653>.
- [18] A. Donatti, S.L. Correa, J.S.B. Martins, A. Abelem, C.B. Both, F. Silva, J.A. Suruagy, R. Pasquini, R. Moreira, K.V. Cardoso, T.C. Carvalho, Survey on machine learning-enabled network slicing: Covering the entire life cycle, *IEEE Trans. Netw. Serv. Manag.* (2023) 1, <http://dx.doi.org/10.1109/TNSM.2023.3287651>.
- [19] R. Pasquini, R. Stadler, Learning end-to-end application QoS from openflow switch statistics, in: *2017 IEEE Conference on Network Softwareization (NetSoft)*, 2017, pp. 1–9, <http://dx.doi.org/10.1109/NETSOFT.2017.8004198>.
- [20] Y. Cui, X. Huang, P. He, D. Wu, R. Wang, Qos guaranteed network slicing orchestration for Internet of Vehicles, *IEEE Internet Things J.* 9 (16) (2022) 15215–15227, <http://dx.doi.org/10.1109/JIOT.2022.3147897>.
- [21] B. Nougnanke, Y. Labit, M. Bruyere, U. Aivodji, S. Ferlin, ML-Based performance modeling in SDN-enabled data center networks, *IEEE Trans. Netw. Serv. Manag.* 20 (1) (2023) 815–829, <http://dx.doi.org/10.1109/TNSM.2022.3197789>.
- [22] Z. Ge, J. Hou, A. Nayak, Forecasting SDN End-to-End latency using graph neural network, in: *2023 International Conference on Information Networking (ICOIN)*, 2023, pp. 293–298, <http://dx.doi.org/10.1109/ICOIN56518.2023.10048915>.
- [23] Z. Ge, J. Hou, A. Nayak, GNN-based End-to-end delay prediction in software defined networking, in: *2022 18th International Conference on Distributed Computing in Sensor Systems (DCOSS)*, 2022, pp. 372–378, <http://dx.doi.org/10.1109/DCOSS54816.2022.00066>.
- [24] F. Laiche, A. Ben Letaifa, I. Elloumi, T. Aguil, When machine learning algorithms meet user engagement parameters to predict video QoE, *Wirel. Pers. Commun.* 116 (3) (2021) 2723–2741, <http://dx.doi.org/10.1007/s11277-020-07818-w>.
- [25] N. Abdelwahed, A.B. Letaifa, S.E. Asmi, Monitoring web QoE based on analysis of client-side measures and user behavior, *Multimedia Tools Appl.* 82 (4) (2023) 6243–6269, <http://dx.doi.org/10.1007/s11042-022-13427-5>.
- [26] S. Khan, S. Khan, Y. Ali, M. Khalid, Z. Ullah, S. Mumtaz, Highly accurate and reliable wireless network slicing in 5th generation networks: A hybrid deep learning approach, *J. Netw. Syst. Manage.* 30 (2) (2022) 29, <http://dx.doi.org/10.1007/s10922-021-09636-2>.
- [27] A. Thantharate, C. Beard, ADAPTIVE6G: Adaptive resource management for network slicing architectures in current 5G and future 6G systems, *J. Netw. Syst. Manage.* 31 (1) (2022) 9, <http://dx.doi.org/10.1007/s10922-022-09693-1>.
- [28] N.P. Tran, O. Delgado, B. Jaumard, F. Bishay, ML KPI prediction in 5G and B5G networks, in: *2023 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*, 2023, pp. 502–507, <http://dx.doi.org/10.1109/EuCNC/6GSummit58263.2023.10188363>.
- [29] H. Yu, Z. Ming, C. Wang, T. Taleb, Network slice mobility for 6G networks by exploiting user and network prediction, in: *ICC 2023 - IEEE International Conference on Communications*, 2023, pp. 4905–4911, <http://dx.doi.org/10.1109/ICC45041.2023.10279739>.
- [30] F. Chiariotti, M. Drago, P. Testolina, M. Lecci, A. Zanella, M. Zorzi, Temporal characterization and prediction of VR traffic: A network slicing use case, *IEEE Trans. Mob. Comput.* (2023) 1–18, <http://dx.doi.org/10.1109/TMC.2023.3282689>.
- [31] R. Dangi, P. Lalwani, Harris hawks optimization based hybrid deep learning model for efficient network slicing in 5G network, *Clust. Comput.* 27 (1) (2024) 395–409, <http://dx.doi.org/10.1007/s10586-022-03960-1>.
- [32] A.C. Bakır, A. Özgövde, C. Ersoy, End-to-end network slicing for edge computing optimization, *Future Gener. Comput. Syst.* 157 (2024) 516–528, <http://dx.doi.org/10.1016/j.future.2024.03.001>, URL <https://www.sciencedirect.com/science/article/pii/S0167739X24000724>.
- [33] G. Barlacchi, M. De Nadai, R. Larcher, A. Casella, C. Chitic, G. Torrisi, F. Antonelli, A. Vespignani, A. Pentland, B. Lepri, A multi-source dataset of urban life in the city of Milan and the Province of Trentino, *Sci. Data* 2 (1) (2015) 150055, <http://dx.doi.org/10.1038/sdata.2015.55>.
- [34] A. Cassandra, Cassandra stress | apache cassandra documentation, 2023, URL [https://cassandra.apache.org/doc/latest/cassandra/tools/cassandra\\_stress.html](https://cassandra.apache.org/doc/latest/cassandra/tools/cassandra_stress.html). (Acessado 04 setembro 2023).
- [35] J. Bergstra, D. Yamins, D. Cox, Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures, in: S. Dasgupta, D. McAllester (Eds.), *Proceedings of the 30th International Conference on Machine Learning*, in: *Proceedings of Machine Learning Research*, vol. 28, PMLR, Atlanta, Georgia, USA, 2013, pp. 115–123.
- [36] G. James, D. Witten, T. Hastie, R. Tibshirani, et al., *An Introduction to Statistical Learning*, vol. 112, Springer, 2013.
- [37] Z. Wang, W. Yan, T. Oates, Time series classification from scratch with deep neural networks: A strong baseline, in: *2017 International Joint Conference on Neural Networks, IJCNN*, 2017, pp. 1578–1585, <http://dx.doi.org/10.1109/IJCNN.2017.7966039>.
- [38] I. Oguiza, Tsai - A state-of-the-art deep learning library for time series and sequential data, 2023, URL <https://github.com/timeseriesAI/tsai>, Github.

- [39] S. Bai, J.Z. Kolter, V. Koltun, An empirical evaluation of generic convolutional and recurrent networks for sequence modeling, 2018, CoRR [abs/1803.01271](https://arxiv.org/abs/1803.01271) [arXiv:1803.01271](https://arxiv.org/abs/1803.01271) URL <http://arxiv.org/abs/1803.01271>.
- [40] H. Ismail Fawaz, B. Lucas, G. Forestier, C. Pelletier, D.F. Schmidt, J. Weber, G.I. Webb, L. Idoumghar, P.-A. Muller, F. Petitjean, InceptionTime: Finding AlexNet for time series classification, *Data Min. Knowl. Discov.* 34 (6) (2020) 1936–1962, [http://dx.doi.org/10.1007/s10618-020-00710-y](https://doi.org/10.1007/s10618-020-00710-y).
- [41] M. Rußwurm, M. Körner, Self-attention for raw optical satellite time series classification, *ISPRS J. Photogramm. Remote Sens.* 169 (2020) 421–435.
- [42] J. Howard, S. Gugger, Fastai: A layered API for deep learning, *Information* 11 (2) (2020) 108.
- [43] T. Akiba, S. Sano, T. Yanase, T. Ohta, M. Koyama, Optuna: A next-generation hyperparameter optimization framework, in: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 2623–2631.
- [44] L. Yang, A. Shami, On hyperparameter optimization of machine learning algorithms: Theory and practice, *Neurocomputing* 415 (2020) 295–316, [http://dx.doi.org/10.1016/j.neucom.2020.07.061](https://doi.org/10.1016/j.neucom.2020.07.061), URL <https://www.sciencedirect.com/science/article/pii/S0925231220311693>.
- [45] C. Mesh, Chaos mesh: Cloud-Native chaos engineering platform, 2024, URL <https://chaos-mesh.org>. (Acessado 11 fev 2025).



**Rodrigo Moreira** received his Ph.D. from the University of Uberlândia (UFU) in 2021. He is currently a Professor at the Federal University of Viçosa, where he also obtained his B.S. degree in 2014. He earned his M.Sc. degree from the Federal University of Uberlândia, Brazil, in 2017. He has several papers published and has presented at conferences. His research interests include the future of the internet, quality of service, cloud computing, network function virtualization, software-defined networking, computational intelligence, and edge computing.



**Rafael Pasquini** received his M.Sc. and Ph.D. in computer engineering from the State University of Campinas in 2006 and 2011, respectively. From 2015 to 2017 he was a Visiting Researcher in the Department of Network and Systems Engineering (NSE) at KTH Royal Institute of Technology. Since 2011, he has been an associate professor and has led the Distributed Systems and Networks (DSN) research group at the Department of Computer Science of the Federal University of Uberlândia. His research interests include network management, slicing of software-defined infrastructures, machine learning, cloud computing, and software-defined networks.



**Joberto S. B. Martins** (Life Senior Member, IEEE) Ph.D. in Computer Science at Université Pierre et Marie Curie (UPMC), Paris (1986), PosDoc at ICSI - Berkeley University (1995), and PosDoc Senior Researcher at Paris Saclay University - France (2016). International Professor at Hochschule für Technik und Wirtschaft des Saarlandes (HTW), Germany (since 2004). Full Professor at Salvador University (UNIFACS) on Computer Science, Director of NUPERC research group with research interests in Network Slicing, Resource Orchestration, Machine Learning, and Smart City. He is a key speaker, teacher, and invited lecturer at various international congresses and companies in Brazil and Europe.



**Tereza Cristina Melo de Brito Carvalho** Associate Professor of Escola Politécnica University of São Paulo (USP) and visiting professor at Université Paris 1 Panthéon-Sorbonne. She has been the founder and general coordinator of LASSU (Laboratory of Sustainability on ITC) since 2010 and CEDIR-USP (Center for Reuse and Discard of Informatics Residuals) since 2009. She is a former assessor of CTI USP (Information Technology Coordination) from 2010–2013 and CCE-USP (Electronic Computing Center) director from 2006–2010. She is a Sloan Fellow in 2002 from MIT (Massachusetts Institute of Technology). She has coordinated international and national R&D projects since 2000 in: Green Computing, Cloud Computing, IT Energy Efficiency, IT Governance, Digital Technologies applied to the Amazon Production Chains, WEEE (Waste Electrical and Electronic Equipment), Future Internet, Scientific DMZ, and Security. She holds several international patents.



**Flávio de Oliveira Silva** (Member, IEEE) received the Ph.D. from the University of São Paulo (USP) in 2013. He is a Professor at the Faculty of Computing (FACOM), Federal University of Uberlândia (UFU). He is also a professor in the Department of Informatics (DI), School of Engineering of the University of Minho in Braga, Portugal, and a researcher with the ALGORITMI Centre. He has published and presented several papers at conferences worldwide. His research interests include future networks, the IoT, network softwareization (SDN and NFV), future intelligent applications and systems, cloud computing, and software-based innovation. He is a member of IEEE, ACM, and SBC. He is a reviewer of several journals and a member of TCP at several IEEE conferences.