

# $\pi$ -ROAD: a Learn-as-You-Go Framework for On-Demand Emergency Slices in V2X Scenarios

Armin Okic\*, Lanfranco Zanzi<sup>†‡</sup>, Vincenzo Sciancalepore<sup>†</sup>, Alessandro Redondi\*, and Xavier Costa-Pérez<sup>†§</sup>

\*Politecnico di Milano, Milano, Italy, Email: {name.surname}@polimi.it

<sup>†</sup>NEC Laboratories Europe, Heidelberg, Germany, Email: {name.surname}@neclab.eu,

<sup>‡</sup> Technische Universität Kaiserslautern, Kaiserslautern, Germany

<sup>§</sup> i2CAT Foundation and ICREA, Barcelona, Spain.

**Abstract**—Vehicle-to-everything (V2X) is expected to become one of the main drivers of 5G business in the near future. Dedicated *network slices* are envisioned to satisfy the stringent requirements of advanced V2X services, such as autonomous driving, aimed at drastically reducing road casualties. However, as V2X services become more mission-critical, new solutions need to be devised to guarantee their successful service delivery even in exceptional situations, e.g. road accidents, congestion, etc. In this context, we propose  $\pi$ -ROAD, a *deep learning* framework to automatically learn regular mobile traffic patterns along roads, detect non-recurring events and classify them by severity level.  $\pi$ -ROAD enables operators to *proactively* instantiate dedicated *Emergency Network Slices (ENS)* as needed while re-dimensioning the existing slices according to their service criticality level. Our framework is validated by means of real mobile network traces collected within 400 km of a highway in Europe and augmented with publicly available information on related road events. Our results show that  $\pi$ -ROAD successfully detects and classifies non-recurring road events and reduces up to 30% the impact of ENS on already running services.

## I. INTRODUCTION

Despite the high investment volume in public transportation, many people daily commute to work with their private vehicles over major roads around cities. Accordingly, both drivers and passengers consume and generate a large amount of data along the road mobile infrastructure for a wide variety of purposes: navigation systems, car sensors, infotainment but also phone calls, social media, streaming, etc. Along highly crowded roads, this may lead to network congestion and, in the worst case, service disruptions.

In order to avoid these situations, the next generation of mobile networks (5G) has defined the *Network Slicing* concept, which allows infrastructure providers to dynamically instantiate *on-demand customized virtual network instances* with dedicated Service Level Agreements (SLAs). Standardization bodies have defined the overarching architecture [1] to support such isolated slices, thus fostering research on dynamic resource orchestration mechanisms based on well-known technologies such as Network Function Virtualization (NFV) and Software Defined Networking (SDN) [2], [3].

As 5G gets rolled-out and advanced V2X services deployed, solutions to protect mission-critical services will become increasingly important. While significant road safety improvements have been introduced in the last decades, road fatalities are still a major cause of injuries and death worldwide [4]. Thus, guaranteeing public safety along roads is still a major

challenge that requires novel solutions. Due to the combined effect of high-mobility patterns, traffic volumes and advanced automotive-related use cases (V2X), mobile network infrastructure along roads will face daunting challenges to guarantee mission critical services in unexpected congestion scenarios. In this context, *Emergency Network Slices (ENS)* are expected not only to improve situation awareness during emergencies but also to support the provisioning of enhanced communication schemes, e.g., Ultra-Reliable and Low-latency Communication (URLLC) for virtual and augmented reality (VR/AR) to first-responder teams, e.g., ambulances, police, firefighters, that have to reach the event location and manage emergencies in a faster and safer manner [5]. As ENS will get top priority, solutions need to be devised to re-dimension already running services according to their criticality level. Unexpected road events progressively cause traffic congestion to nearby areas and, in turn, saturate the networking resources of adjacent base stations. If such a *propagation effect* could be *predicted*, their congestion effects could be alleviated by means of *proactive slices resource management*.

We take on this challenge and propose *Passive Information-based Resource Orchestration in Automotive Driving scenarios* ( $\pi$ -ROAD) that relies on a deep learning framework to *analyze* passive information from real-time mobile network traffic statistics, *learn* regular traffic patterns, and accurately *detect* anomalous deviations due to unexpected road events.

Our contributions can be summarized as follows: *i)* we provide an in-depth analysis of *real mobile network traces* and statistics from an operational network along 400 Km of a major road in Europe, *ii)* we design a deep learning model, namely  $\pi$ -ROAD, that applies the well-known concept of learn-as-you-go, i.e., it accurately detects and localizes road events and classifies them according to their severity, *iii)* we focus on the network orchestration of emergency scenarios where an Emergency Network Slice (ENS) must be instantiated according to the type of road event, *iv)* we formulate an optimization problem to minimize the impact of ENSs on already existing slices and *v)* we validate our model and assess its capabilities by means of a full dataset of network monitoring data in realistic scenarios.

The remainder of the paper is structured as follows. Section II provides an overview of road dynamics from a mobile network perspective. Section III details the main building

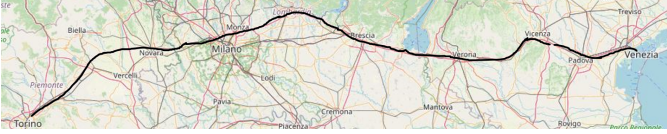


Figure 1: Overview of the Italian highway (A4).

blocks of our solution. Leveraging on the output of  $\pi$ -ROAD, Section IV formulates the ENS problem, whereas Section V presents an exhaustive simulation campaign to validate our design principles through real operational data. Section VI summarizes the related literature and VII concludes the paper.

*Privacy issues:* The research activity presented in this paper does not violate user's privacy rights. The dataset contains only aggregated and anonymous information collected by the network operator for monitoring purposes.

## II. DATA ANALYSIS

We carry out an analysis of the network traffic dynamics occurring on the vehicular roads.

The considered dataset consists of 6 months of real network data (February-July 2020) collected from an operational LTE network deployed alongside the Italian A4 highway shown in Fig. 1. The highway has a length of approximately 400 km and is located in the north of Italy, connecting the city of Turin with Venice, passing through the metropolitan area of Milan. Along this highway segment, more than 1000 LTE cells, corresponding to about 200 eNodeBs (eNBs), are deployed to provide connectivity to the users commuting or traveling on this path, and to citizens leaving in the proximity of the highway. Available data exploit local monitoring information of LTE eNBs including both averaged and aggregated Key Performance Indicators (KPIs) within a time granularity of 15 minutes.

### A. Spatial consideration

Fig. 2 (top) depicts a snapshot of normalized DL and UL traffic volumes (per single base station, where the approximate base station location is highlighted in the upper part of the plot) aggregated over a month. We highlight areas of the main cities located nearby the A4 roadway. As expected, DL volumes increase in the proximity of major cities, due to urban traffic leakage of base stations covering the highway. Additional traffic peaks can be identified in specific geographical locations, e.g., intersections between different highways, main inter-urban roads, or train lines. Generally, these locations are characterized by a higher density of base stations, which means that occurrence of any road events in these points will probably lead to major impact on the mobile network. Moreover, a large number of mobile terminals characterizes such locations over time, due to a significantly higher user mobility. Similarly, UL traffic shows analogous behaviors but with a limited volume (about 10% of the overall DL traffic).

### B. Temporal patterns consideration

Several works in the literature suggest a strong relationship that correlates end-users mobility patterns with cellular network statistics [6]–[8] in urban environments. In a similar way,

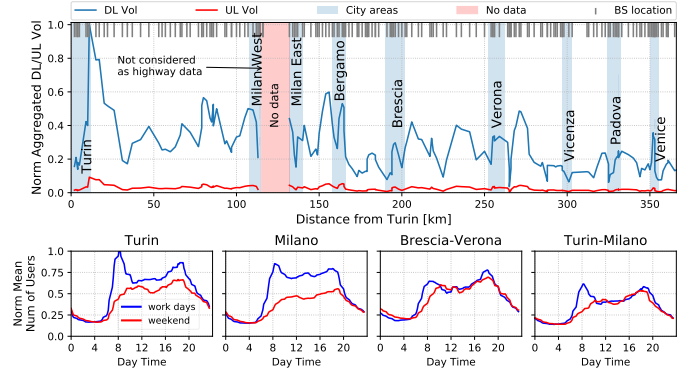


Figure 2: Aggregated spatial distribution of DL/UL traffic volumes along the highway (top), with the focus on working day and weekend signatures for specific highway areas (bottom).

the mobility patterns identified on highways present repetitive trends following regular working routines. In this respect, we plot at the bottom of Fig. 2 normalized number of active users during working days and weekends for different areas of the highway, in particular two areas around metropolitan cities as well as two sections of the highway between two main cities. The signatures are extracted by calculating the median values of same time periods over several weeks, separately for working days of the week and weekends, as proposed by [9].

The first and second subplots show that areas around big industrial cities, such as Milan and Turin, are characterized by commuting patterns with a presence of mobility peaks in the morning, noon and evening during the working days, which are absent on weekends. Conversely, mobility patterns of users during working and weekend days are comparable when considering highway segments interconnecting major cities (third and fourth subplots).

### C. Service-based QoS consideration

Different services with heterogeneous requirements might be seamlessly managed by the mobile network infrastructure. In real deployments, this is usually achieved by labelling each traffic flow with a specific QoS Class Identifier (QCI) to ensure that each traffic bearer [10] is allocated with the appropriate set of resources to guarantee an affordable Quality of Service (QoS). In Fig. 3 (left-hand side), we depict the temporal distribution of the DL volumes differentiated by QCI traffic types. It can be noticed a dominance of non-guaranteed bit-rate (NON-GBR) traffic types (QCIs 6 to 9), mostly related to video-streaming and social media activities, and almost negligible volumes for guaranteed bit-rate (GBR) traffic types (QCIs 1 to 4). Note that the satisfaction of service requirements also depends on the network congestion level and on the instantaneous channel quality experienced by end-users, together with the corresponding modulation and coding scheme (MCS) selected at the eNodeB scheduler. Interestingly, DL traffic perceives, on average, a lower MCS index with respect to the UL one (right-hand side of Fig. 3). This is due to the high end-user mobility and longer data exchange sessions, which suffer from a wider communication distance along the path [11].

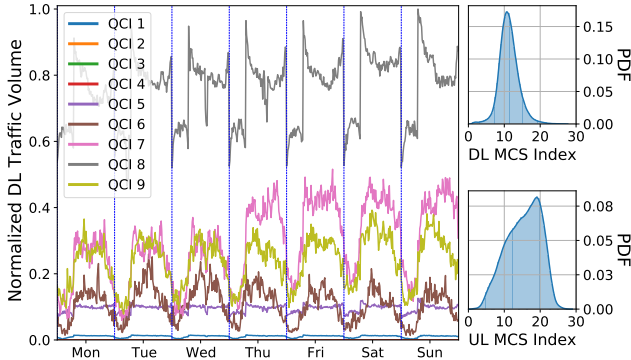


Figure 3: Temporal distribution of different traffic types volumes.

#### D. Road events consideration

The considered dataset takes into account monitoring information from the mobile network infrastructure, but it lacks of information related to vehicular traffic and unexpected congestions/events that affect its dynamic. To fill this gap and help us to accurately relate network traffic dynamics with road events, we rely on publicly available information coming from heterogeneous social media sources like Twitter notification service of *Autostrade per l'Italia*<sup>1</sup>, Google Maps<sup>2</sup> and Waze<sup>3</sup>, which facilitate an exhaustive catalog of real-time road event information and useful metadata, e.g., timestamp, exact location and a short description.

After 6 months of data collection, the final dataset includes about 800 road events,<sup>4</sup> as shown in Fig. 4. As expected, road events mainly occur during morning and evening commuting periods, and are geographically placed close to main cities, e.g., Milan and Verona. We will match the information contained in this supplementary dataset with the network geographical deployment information to assign each road event with the closest base station and obtain the ground truth of occurred events, as later detailed in Section III.

#### E. Event propagation effect

Road events might have a consistent impact on the vehicular traffic conditions and, in turn, on overall network resource availability. However, it might be easily confused with common network traffic outliers and therefore be ignored. Current network deployments may come to help: the RAN deployment along the highway has quite a regular pattern and might reveal implicit information as monitoring data are simultaneously retrieved from different observation points. In particular, a road event might affect the end-users activity or drop the mobility rate gradually for adjacent base stations. We call it *event propagation effect*. It usually depends on the severity of the road event, the day time, the specific location and the network deployment (e.g., base station density, etc.).

<sup>1</sup>Online available at: <https://twitter.com/trafficoa>

<sup>2</sup>Online available at: <https://www.google.com/maps>

<sup>3</sup>Online available at: <https://www.waze.com/livemap>

<sup>4</sup>Note that our experimental campaign has been carried out during the Covid-19 pandemic, when the imposed lockdown limited the overall end-user mobility in the north of Italy, therefore, reported numbers may be biased and differ from yearly regular data.

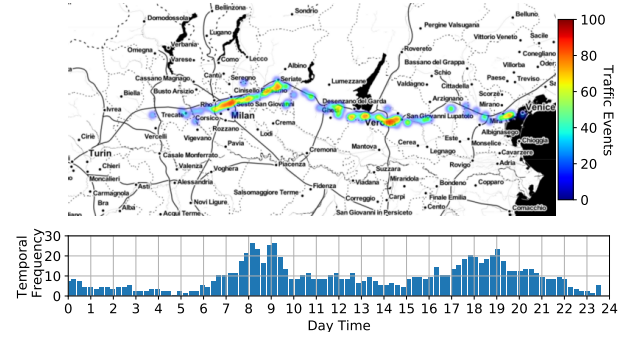


Figure 4: Spatio-Temporal distribution of road events.

From the auxiliary information contained in our event dataset, we make a straightforward association between the location of each event and the closest base station, thereby identifying the source of the road event in the network. Then, analyzing the propagation of anomalous traffic patterns onto nearby cells, we can infer which direction of the highway, or vehicular traffic flow, has been affected.

An example of this is shown in Fig. 5. The scenario accounts for a major accident occurred nearby Verona at 9:00 AM. The red line in each plot represents the regular behaviour of the corresponding metric, while the blue line depicts the daily data. After the accident, we can notice a clear deviation from regular patterns for multiple metrics lasting few hours. The higher number of users in RRC connected state indicates that multiple users are actively using mobile network resources, suggesting the presence of traffic congestions. Similarly, this affects DL volumes and average Physical Resource Block (PRB) utilization. Moreover, a higher number of affected base stations on the east side suggests that the accident occurred in east-west direction (Venice-Turin). Given the high variability and intrinsic characteristic of each monitoring metric, an accurate anomalous pattern characterization would involve time, space and metric-specific considerations. Therefore, road events can hardly be detected by simple anomaly detection algorithms, e.g., those based on outlier detection that may

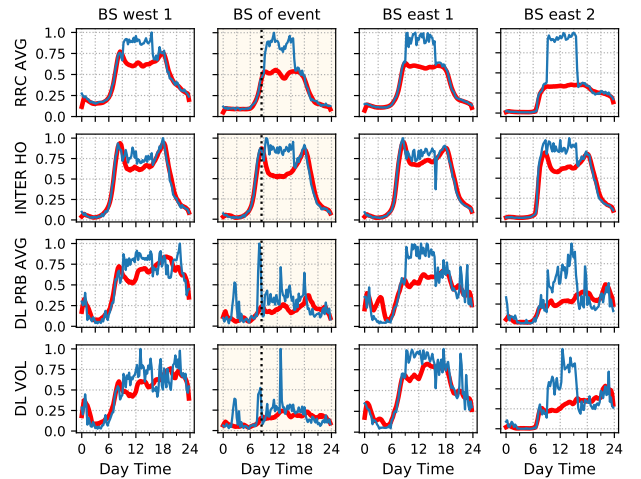


Figure 5: Event propagation on nearby base stations in both directions for selected metrics.

run on a single base station. This motivates us to investigate advanced deep learning solutions capable of dealing with such multi-dimensional information matrix while keeping a global view of the system.

### III. $\pi$ -ROAD DESIGN AND MODEL FORMULATION

Hereafter, we introduce the design of  $\pi$ -ROAD, a deep learning-based framework to detect road events. An overview of the model architecture is depicted in Fig. 6. The overall solution requires two different stages to deal with road event predictions: *i*) an autoencoder-based approach consisting of Long Short-Term Memory Recurrent Neural Networks (LSTM) layers to detect anomalies within highly dynamic and heterogeneous time patterns and *ii*) a Deep-Learning approach to approximate complex functions, e.g., those that relate road events occurrences with temporal and spatial distributions of anomalies detected in adjacent base stations. The overall 6 months dataset has been splitted according to a 60/20/20 ratio for the purposes of training, validation and testing procedures, respectively.

#### A. Input of $\pi$ -ROAD

Let us consider a time slotted system whereby time is divided into time intervals  $t = \{1, \dots, T\}$ , and define  $\mathcal{N}$  as the set of base stations deployed along the highway. As detailed in Section II, each base station  $n \in \mathcal{N}$  gathers information about a multitude of heterogeneous network statistics, defined as  $\mathcal{M}$ , which include, among the others, PRB utilization, physical channel quality information, traffic volumes, X2/S1 hand-over statistics, etc. In our experiments we consider 25 different monitoring metrics.

Let  $x_m^n(t)$  be the time series describing the  $m$ -th monitoring metric observed under the base station  $n$  at the time  $t$ . As suggested in [12], an accurate organization of the input traces helps enhancing deep learning models performances. In light of this, for all metrics  $m \in \mathcal{M}$  and base stations  $n \in \mathcal{N}$ , we collect the  $x_m^n(t)$  traces and order them by preserving their spatial location (from Turin to Venice) in a matrix  $X_m(t)$  with dimensions  $N \times L$ , representing a metric-specific *snapshot* of the network, where  $N$  is the total number of base stations, and  $L$  is the number of monitoring samples within the observation period (or *lookback* time window), fixed to 4 hours throughout our experiments. The input matrices  $X_m(t)$  are normalized with respect to maximum values of each  $m$  in the training dataset.

#### B. LSTM-based anomaly detection

The possibility to detect changes in traffic conditions is subject to the capabilities of our model to correctly identify irregular statistics from the monitoring samples collected along the highway. Given the multi-dimensionality of our dataset, we build an autoencoder  $A_m$  for every collected metric thereby allowing for better scalability. Autoencoders imply the setup of an encoding-decoding architecture. The encoding part, commonly implemented as feed-forward neural networks, provides a compressed representation of the input data to subsequent layers [13]. The decoding phase follows the same steps over

a symmetric architecture. Analytically, the two phases applied to the input metric  $X_m$  can be described as follows:

$$h_m(X_m(t)) = \eta(W_m^\eta X_m(t) + b_m^\eta), \quad (1)$$

$$\hat{X}_m(t) = \delta(W_m^\delta h_m(X_m(t)) + b_m^\delta), \quad (2)$$

where  $\eta$  and  $\delta$  are the encoding and decoding functions with their corresponding weights  $W_m^\eta$  and  $W_m^\delta$ ,  $b_m^\eta$  and  $b_m^\delta$  are bias vectors, and  $h_m(X_m(t))$  and  $\hat{X}_m(t)$  are the compressed input and the reconstructed output sequences, respectively.

The performance of autoencoders, defined as their capability to reconstruct the input sequence, depends on variability and complexity of data provided at the input. Differently, in our work we exploit autoencoders for anomaly detection [14] instead of simple dimensionality reduction.

The data analysis performed in Section II suggests the adoption of deep learning techniques able to deal with the spatio-temporal characteristics of mobile traces. For this reason, we implement our encoder-decoder architecture by means of Long Short-Term Memory Recurrent Neural Networks (LSTM-RNNs). Each autoencoder  $A_m$  accounts for 4 LSTM layers, two for each phase. LSTM is a type of RNN architecture that has proven its value when dealing with repetitive patterns and unstructured time series, while solving a problem of vanishing gradient for long-term dependencies present in other RNNs [15]. LSTMs can be represented as a chain of  $G$  modules, or cells, each one applying a set of operations to the input data. In our case, the two LSTM layers are composed by 128 and 64 cells, respectively [16].

The output of each cell  $g \in G$ , also known as cell state, is transferred to the subsequent cell in a recursive manner. The possibility of handling long-term trends in LSTM is provided by structures, dubbed as gates, which carefully remove or add information to the cell state. Each cell has three gates, namely, input  $I_g$ , output  $O_g$  and forget gate  $F_g$ , which controls the amount of information that should be added (or dropped) to the cell state before transferring it to the next unit [17]. This effect is achieved combining the influence of different non-linear activation functions, i.e.,  $\sigma$  (sigmoid) and  $\tanh$  (hyperbolic tangent function), at each gate. The impact of these functions on the input data needs to be learned during the training phase aiming to minimize a loss function [17], which in our case is the mean squared error (MSE).

To accomplish the anomaly detection task, we make use of labeled data and train the autoencoders offline exclusively on *eventless* snapshots taken from historical data as to capture the behavior of the system without anomalies. Once trained, we feed the model with online monitoring traces. Given that anomalous patterns have not been part of the training phase, it is expected that the model will exhibit performance degradation during the reconstruction phase. Simple classifiers would mark the input snapshot as anomalous if the reconstruction error exceeds a given threshold. Instead, we derive the squared error matrix  $e_m(t) = \|X_m(t) - \hat{X}_m(t)\|^2$  from each autoencoder  $A_m$  and combine the individual error matrices into a 3D tensor  $e(t) = \{e_1(t), e_m(t), \dots, e_M(t)\}$  which is passed to the second stage of our model.

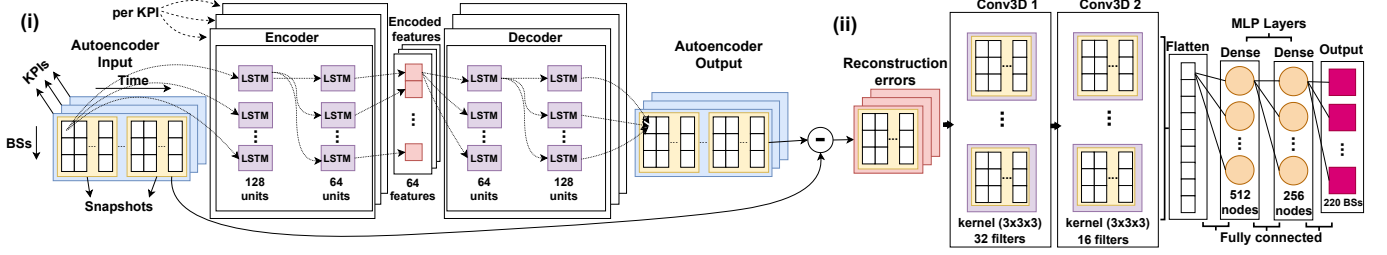


Figure 6:  $\pi$ -ROAD model architecture

### C. Road Event Localization

The function linking anomalies detected by the first stage of our model with the occurrence of road events along the highway is unknown and hard to be characterized due to the multitude of system-related and external variables affecting the entire detection process. The simple detection of anomalies in monitoring statistics does not imply the occurrence of an emergency, as uncorrelated events—like hardware failures and/or unexpected traffic peaks—may trigger alerts leading to erroneous estimations. For this reason in the following, we exploit the information described in Section II-D as ground-truth to train a Convolutional Neural Network (CNN)-based model that captures the spatio-temporal correlation of different anomalies and maps them into geographical information [18]. Our design choice is further motivated by the fact that in case of road events, as shown in Fig. 5, affected base stations present significant levels of correlation between each other, which further improves the learning task.

More in details, the second stage of  $\pi$ -ROAD consists of two stacked 3D-CNN layers and a final Multi-Layer Perceptron (MLP) fully connected layer. In order to exploit local correlation from nearby base stations, each neuron of the 3D-CNN layers has a limited *receptive field*, or kernel, whose size determines its observation area. For a given input tensor, convolutions with 3D kernels are iteratively applied to provide the subsequent layers with a compressed representation of the input information.

Through extensive hyper-parameters optimization, we used two 3D-CNN layers with filter sizes of 32 and 16 neurons, respectively, and corresponding convolutional kernel sizes of (3, 3, 3) and (3, 3, 3) [19]. Each neuron of the CNN runs a filter  $H(\sum_t e(t) * k(t) + b)$ , where  $e(t)$  is the input tensor at time  $t$ ,  $k(t)$  is the kernel filter,  $*$  is the convolution operator,  $b$  is a bias, and  $H(\cdot)$  is a non-linear activation function, in our case Rectified Linear Unit (ReLU) [20].

In order to map the encoded representation of the anomalies into geographical information about the emergency, we make use of a final Multi-Layer Perceptron layer. MLP is a class of neural networks with fully-connected neurons among layers which has the capabilities to approximate, through supervised learning, the function that relates the input with the output. In our case, we are interested in the function that links the encoded spatio-temporal anomalies of multiple metrics and different base stations with the exact event location. The MLP classifier consists of three layers with 512, 256 and 220 neurons, respectively, where the output layer matches

the number of base stations [21]. To regularize the output and reduce over-fitting, we introduce a dropout rate of 0.2. Overall, the second stage of the model is trained using Adam optimizer with learning rate  $10^{-4}$ , and adopting cross-entropy as loss function over 150 training epochs. Due to high system variability caused by user social behaviors and external causes affecting it (e.g., weather conditions), it clearly rises the need to design a model capable of adapting to new (anomalous) patterns. Therefore, we ensure that the model is retrained as soon as new observations are made available.

### D. Road Event Classification

Upon detection, it is important to classify the magnitude of each event to promptly alert first aid responders and, if necessary, identify the set of networking requirements to be allocated for an emergency slice setup. We rely on the impact of the *propagation effect* over the set of base stations near to the road event to provide an empirical classification. Let us introduce  $\tilde{\mathcal{N}} \subset \mathcal{N}$  as the set of base stations affected by the road event. Due to the irregular density that characterizes the Radio Access Network (RAN) deployment along the highway, we argue that a simple road event classification based on the cardinality of  $\tilde{\mathcal{N}}$  would not be accurate, as road events occurring in the proximity of main cities would involve a wider number of base stations than those occurring in rural areas. Therefore, we define our classification metric  $\mu$  as  $\frac{\tilde{N}}{\psi}$ , where  $\tilde{N}$  is the cardinality of  $\tilde{\mathcal{N}}$ , and  $\psi$  is the number of base stations deployed in the area within a radius of 10 km surrounding the road event. We empirically select 10 km as this value represents the longest vehicular queue registered in our event dataset. Clearly,  $\psi$  should be re-dimensioned to generalize our findings within other network topologies.

Finally, we differentiate among three different categories of events: *Light*, *Moderate*, and *Severe*. The event duration of each category reflects the most common scenarios contained in our real dataset, while networking throughput requirements are meant to support the provisioning of Augmented and Virtual Reality (AR/VR) streaming in mission critical scenarios [22]. The details about each road event class and corresponding ENS requirements are summarized in Table I.

Table I: Event characterization and ENS Requirements.

Event category	$\Theta$ - Time Duration	NS Requirements	$\mu$ Value
Light	20min	10Mbps	(0, 1]
Moderate	40min	15Mbps	(1, 2]
Severe	60min	25Mbps	$\geq 2$

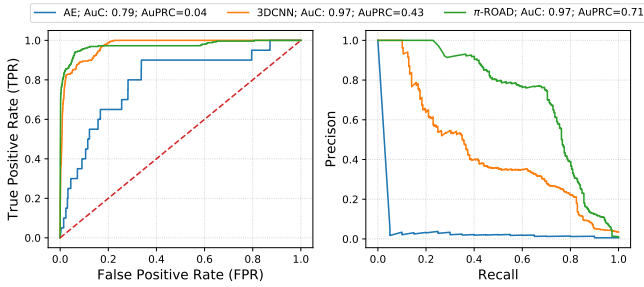


Figure 7: ROC and PR diagrams with indicated AuC and AuPRC scores for different anomaly detection and classification approaches.

#### E. Performance Comparison and Practical Considerations

Hereafter, we compare the performance of  $\pi$ -ROAD against state-of-the-art models trained to perform similar anomaly detection and classification tasks, highlighting each model drawback (compared to our solution) through practical considerations. Benchmarks include a simple threshold-based autoencoder classifier (AE) [16], and a more advanced 3D-CNN-based classifier [19], [21]. Fig. 7 resumes the overall performances using Receiver Operating Characteristics (ROC) and Precision-Recall (PR) metrics. On the left-hand side, the ROC metric provides a compact representation of the capabilities of the model to deal with the trade-off between True Positive Rate (TPR) and False Positive Rate (FPR). On the right-hand side, the PR curve states the performance of a classifier in terms of Precision and Recall, where Precision is a measure of result relevancy, and Recall is a measure of how many truly relevant results are returned by the classifier. To ease the comparison, we also quantify the Area under the ROC Curve (AuC) and Area under the PR Curve (AuPRC). The unsupervised approach adopted in AE requires the definition of a-priori thresholds to mark anomalous patterns from the reconstruction errors. Clearly, the definition of these settings depends on the intrinsic variability of the traces in input, and requires considerations over the statistics of each feature. This approach hardly scales when considering multiple heterogeneous metrics, as in our case. Despite accurate tuning of the parameters, it clearly appears how this simple approach fails to reveal the majority of events, achieving the lowest precision score. This further emphasizes how detection schemes based on simple threshold/outlier detection are not enough to address the scenario considered in our work. Conversely, the 3D-CNN classifier adopts a supervised approach, which prevents from specifying user-defined thresholds for decision-making. The resulting ROC and PR curves show better performances when compared against the baseline AE approach. This result can be explained through the ability of 3D-CNN network to capture spatio-temporal correlations between different network measurements [23]. However, when Recall value increases over a certain level (i.e., moving from left to right on the PR diagram), the Precision score drops drastically, suggesting poor performances when differentiating among different types of anomalies. In other words, the model detects only some types of anomalies, e.g., those deriving from major events with

very strong impact on monitoring traces, and fails to generalize over probably smaller ones. Conversely,  $\pi$ -ROAD outperforms the two stand-alone approaches in both ROC and PR metrics. The advantage deriving from the two-stage approach adopted by  $\pi$ -ROAD is two-fold *i*) the initial anomaly detection task performed by first stage of the model, together with an accurate input organization, facilitates the learning task of the 3D-CNN layer and *ii*) the capabilities of the 3D-CNNs to deal with bare reconstruction errors removes the need to provision user-defined thresholds which may bias the final results.

#### IV. ENS ORCHESTRATION

The outcome of  $\pi$ -ROAD can be used to tackle a variety of public safety issues along the highway. For example, in case of road accidents, it is important to enable dynamic RAN resource allocation in order to provide first responder teams with enough communication capabilities, regardless of active user sessions or services. To this aim, in what follows we introduce our formulation to address the Emergency Network Slice (ENS) orchestration problem.

**Radio Access Network.** Let us consider a RAN deployment with network slicing support covering the area surrounding the highway and comprising a set of base stations  $\mathcal{N}$ . Each base station  $n \in \mathcal{N}$  is characterized by a capacity  $C_n$ , defined as the sum of its available physical resource blocks (PRBs). As result of the road event localization process executed by  $\pi$ -ROAD and described in Section III-C, we identify the subset of base stations  $\tilde{\mathcal{N}} \subset \mathcal{N}$  that will host the emergency slice.

**Active slice services.** Let us assume a set of network slices  $\mathcal{S}$  supporting V2X services, being already installed and active on the considered RAN deployment, whereas let us mark  $s = 0$  the Emergency Network Slice (ENS) to be temporarily installed. We assume each V2X slice  $s \in \mathcal{S}$  described by means of a predefined network slice template that suggests slice requirements in terms of throughput  $\Delta_s$ , and latency  $\Lambda_s$ . Typical values for V2X services are shown in Table II.

##### A. Problem design

In real scenarios, when setting up an emergency slice, advanced orchestration operations are required to still guarantee service level agreements (SLAs) of active slices. Assuming that an admission control process has been executed to accept and install V2X slices on the network with the aim of maximizing the resource efficiency (while still honouring expected service requirements), it might appear challenging to add on top of active slices an additional high-priority service, such as ENS. However, an efficient admission control will accommodate slices with heterogeneous requirements to compensate unexpected slice behaviors or traffic peaks [27]. Analytically, we assume that predefined SLAs include for each slice  $s \in \mathcal{S}$  and base station  $n \in \mathcal{N}$  a minimum number of

Table II: V2X Slice requirements (c.f. [24]–[26])

V2X category	Latency	Data rate	Reliability
Autonomous driving	10ms	10Mbps	99.999%
Tele-operated driving	20ms	25Mbps Sensor Data Streaming	99.999%
Vehicular Internet/Infotainment	> 100ms	15Mbps Video Streaming	Not Specified
Road Safety	100ms	1Mbps	99%

reserved PRBs, namely  $Q_{n,s}^{(t)}$ <sup>5</sup>. Without loss of generality, we assume that the overall resource availability is assigned to the set of running slices. In addition, we translate the slice latency requirements  $\Lambda_s$  into a tolerance value  $\lambda_s$  that defines the maximum number of time intervals packets shall wait into the queue before being served, or dropped due to lack of resource availability within the latency requirements. Note that if all running V2X slices have high priority, i.e., very low delay tolerance, it might be infeasible to install the ENS without impairing other active services. ENS parameters depend on the severity of the emergency road event which can be obtained as output of a  $\pi$ -ROAD execution. We resume the corresponding settings in Table I, in light of the discussions of Section III-D. We define  $\Theta$  as the envisioned emergency time duration, and assume, for each ENS to be deployed, a fixed amount of PRBs allocation  $Q_{n,0}^{(t)}, \forall t \in \{0, \dots, \Theta\}$  and the lowest delay tolerance  $\lambda_0 = 1$ . If not properly scheduled, the additional ENS may lead in the worst case to resource deficit and service disruption in one or multiple base stations. Therefore, we aim at minimizing the instantaneous resource deficit  $\pi_n^{(t)}$ , while still guaranteeing defined SLAs for other active running slices. Our problem can be formulated as follows:

**Problem 1** ( $\pi$ -Orchestrator).

$$\begin{aligned}
& \text{minimize} && \sum_{n \in \tilde{\mathcal{N}}} \sum_{t \in \mathcal{T}} \pi_n^{(t)} \\
& \text{subject to} && \sum_s z_{n,s}^{(t)} \leq C_n + \pi_n^{(t)}, \quad \forall n \in \tilde{\mathcal{N}}, t \in \mathcal{T}; \\
& && \sum_{t=0}^{\Theta} Q_{n,s}^{(t)} - z_{n,s}^{(t)} \leq \sum_{t=\Theta+1}^{\Theta+\lambda_s} Q_{n,s}^{(t)}, \quad \forall s \in \mathcal{S}, n \in \tilde{\mathcal{N}}; \\
& && z_{n,0}^{(t)} = Q_{n,0}^{(t)}, \quad \forall n \in \tilde{\mathcal{N}}, t \in \mathcal{T}; \\
& && z_{n,s}^{(t)} \in \mathbb{N}^+, \quad \forall n \in \tilde{\mathcal{N}}, s \in \mathcal{S}, t \in \mathcal{T}; \\
& && \pi_n^{(t)} \in \mathbb{N}^+, \quad \forall n \in \tilde{\mathcal{N}}, t \in \mathcal{T}.
\end{aligned}$$

where (integer) decision variables are  $\pi_n^{(t)}$  and  $z_{n,s}^{(t)}$  indicating the number of PRBs to be assigned to slice  $s$  on base station  $n$  at time interval  $t$ . Note that  $z_{n,0}^{(t)} = Q_{n,0}^{(t)}$  is due to the highest priority assigned to the ENS ( $s = 0$ ). The first set of constraints introduces flexibility into the problem by adding a non-zero fictitious value (i.e., the resource deficit  $\pi_n$ ) to avoid infeasible solutions. The second set of constraints specify that the slice resource reservation must be performed fulfilling the slice delay tolerance, i.e., waiting traffic still in the queue must be scheduled within  $\lambda_s$  time slots. We run Problem 1 for a set of time slots  $\mathcal{T}$  that includes the time window  $\Theta$  required by the ENS. Problem 1 is an Integer Linear Programming (ILP), which can be efficiently approximated by means of relaxation techniques (e.g., [30]), and commercial solvers to provide near-optimal results<sup>6</sup>.

<sup>5</sup>When SLAs disclose information on expected throughput or user rate, the amount of required PRBs could be dynamically obtained by monitoring the traffic demand of the different slices. For further details, we refer the reader to [28], [29].

<sup>6</sup>The implementation of the problem has been carried out using the framework of IBM ILOG CPLEX and its Python API.

**Proposition 1.** *Problem 1 is NP-Hard and difficult to approximate within  $N^{1-\epsilon}$ , for  $\epsilon > 0$ .*

*Proof:* The NP-Hardness proof goes by reduction. Problem 1 clearly belongs to the NP problems. Let us consider an instance of a bin packing problem (BPP) with arbitrary constants  $a$  and  $B$ , and binary decision variables  $x_{ij}$  and  $y_j$ , where  $\mathcal{I}$  and  $\mathcal{J}$  are the sets of items and bins, respectively [31].

**Problem 2** (BPP).

$$\begin{aligned}
& \text{minimize} && \sum_j y_j \\
& \text{subject to} && \sum_i a x_{ij} \leq B y_j, \quad \forall j \in \mathcal{J}; \\
& && \sum_j x_{ij} \leq 1, \quad \forall i \in \mathcal{I}.
\end{aligned}$$

If we consider  $N = 1$  base station, and  $z_{n,s}, \pi_n$  as binary decision variables, Problem 1 reduces to a BPP that shows that for all  $\epsilon > 0$ , packing items within the minimum number of bins within  $N^{1-\epsilon}$  is NP-hard [31]. Since such a problem is trivial compared to our original Problem 1—that includes a number of base stations  $N > 1$  and integer decisions variables—this result is rather strong. ■

#### B. V2X traffic scheduling

The instantaneous resource deficit  $\sum_{n \in \tilde{\mathcal{N}}} \pi_n^{(t)}$  depends on the severity of the detected road event and, in turn, on the specific ENS resource reservation and time duration settings. Therefore, Problem 1 considers the worst case scenario, i.e., when each active slice fully utilizes reserved resources. However, as shown in Section II, network congestions rarely occur outside the commuting time periods and some of reserved network resources may be underutilized. So, the impact of the ENS on the system can be further mitigated relaxing the fixed PRB allocation scheme envisioned to assure slice isolation at RAN scheduling level [32], [33]. Specifically, if a network slice is underutilized, other slices with pending traffic can use some of its network resources (PRBs)<sup>7</sup>. This allows us to devise a practical algorithm that reduces the overall slice resource deficit. In particular, the algorithm sorts slice traffic requests based on the slice priority (or traffic delay tolerance). Slices with higher priority will use assigned PRBs as per the solution of Problem 1 ( $z_{n,s}^{(t)}$ ) within time interval  $t$  to serve incoming traffic requests. Once all requests have been served, remaining slice PRBs ( $z_{n,s}^{(t)} > 0$ ) are used to serve traffic requests of the next slice in the priority-ranked list. When all PRBs are used ( $\sum_s z_{n,s}^{(t)} = 0$ ), the algorithm proceeds to the next time interval  $t$  keeping unserved traffic requests in the slice queues. If slice traffic requests are not served within the slice delay tolerance  $\lambda_s$ , these are dropped. The resulting resource deficit is further minimized as the RAN slice scheduler efficiently assign resources to slices with pending traffic. Results are shown and discussed in Section V-B.

<sup>7</sup>While this concept may fail to comply with the network slice isolation principle, we assume that the isolation is still guaranteed at higher (abstracted) scheduling layers [32].

## V. $\pi$ -ROAD PERFORMANCE EVALUATION

In the following, we investigate the performance of  $\pi$ -ROAD while detecting road events, and use this information to assess our solution's capabilities in orchestrating radio resources in case of road accidents and realistic traffic conditions involving different types of V2X network slices.

### A. Localization of the road event

In order to evaluate event localization performances, we provide as input the test set of monitoring samples taken from our dataset. In Fig. 8, we summarize the results, focusing on the set of base stations and a representative time period. Each point on the map represents a pair of base station and time interval, and its color indicates the detected probability of event occurrence given the monitoring information. Circular markers on the same map highlight the ground-truth location of road accidents. From the picture it can be noticed how  $\pi$ -ROAD detects most of the events, not only revealing their temporal duration (x-axis), but also suggesting how much the propagation effect influences adjacent base stations (y-axis). Interestingly, the heatmap also reveals some cases in which  $\pi$ -ROAD detects the occurrence of events before online notification services and social media platforms. For visualization purposes, at the top of the figure we highlight two geographical areas of the highway covering impacted base stations in two different event occurrences happening at time  $t_{77}$  and  $t_{84}$ , respectively. In the first case,  $\pi$ -ROAD classifies the road event as *Light* since only two base stations reported anomalous statistics. In the second case, given the wider number of affected eNodeBs, the event is classified as *Moderate*. Finally, few events are not detected. Either this can be due to very limited impact of such events on the monitoring traces or, they may be approximated with patterns that are unknown to the model. In the latter case, we expect that similar events would be detected in the future as the model follows the learn-as-you-go approach to continuously train with the latest data.

### B. Resource Orchestration

Hereafter, we test our ENS orchestration solution exploiting  $\pi$ -ROAD's outputs and the realistic monitoring information contained in our dataset. We consider a 4-slices network scenario, including two enhanced Mobile Broadband (eMBBs) slices dedicated to streaming and infotainment services (with high delay tolerance), and two Ultra-reliable low-latency communication (URLLC) slices dedicated to Autonomous and Tele-operated driving, with corresponding networking requirements as detailed in Table II. We assume that DL slice traffic volumes are generated under the reference set of base stations in the considered event period, according to the specific service QCI and corresponding traces as depicted in Fig. 3. In particular, eMBB slices are mapped to QCIs traces 7 and 8, while URLLCs to QCIs 1 and 6, respectively.

Intuitively, the impact of an ENS installation over the radio access network depends, among the others, by the networking capacity deployed in the area of the road event, instantaneous traffic demands, slice requirements, and user density. In Fig. 9a (top) we provide a quantitative measure in case of different

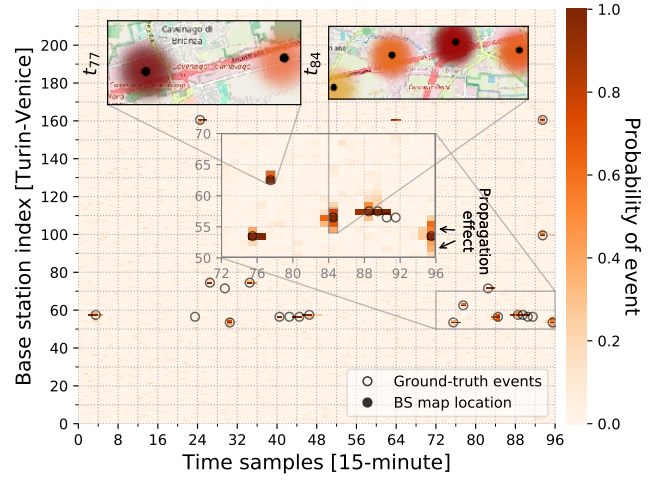
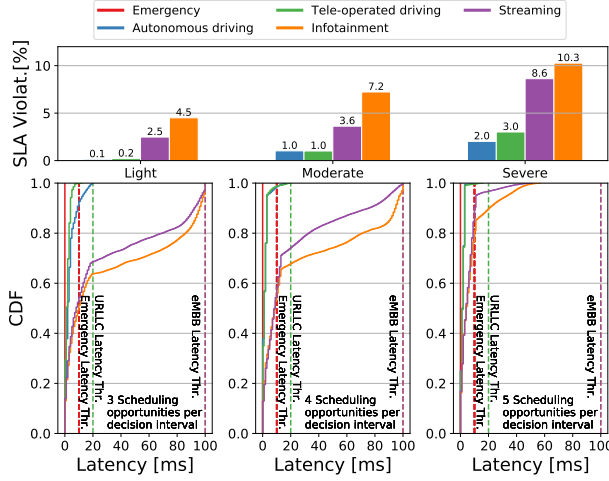


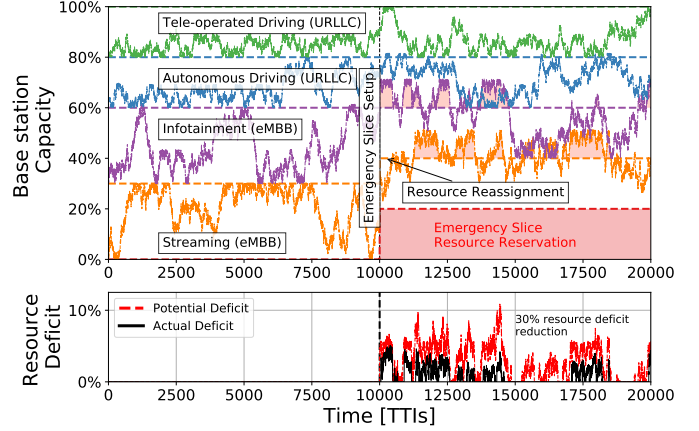
Figure 8: Example of road event detection and localization using  $\pi$ -ROAD model along the highway within one day, compared against the event ground-truth information.

road event types and corresponding duration. We consider channel conditions as perceived by the base station during the event, thereby limiting the base station capacity (due to lower assigned MCS) when compared to the ideal scenario. The results of each category are averaged over 10 road event occurrences taken from our real dataset. It can be noticed that eMBB slices unveil worst performance (w.r.t. URLLC) due to their higher throughput requirements and lower slice priority (i.e., higher delay tolerance). However, an advanced slice orchestration solution may alleviate this issue.

Fig. 9a (bottom) focuses on the *Moderate* road events evaluating the impact of different slice types (delay tolerance parameters  $\lambda_s$ ) with a fixed decision interval  $t$  duration of 100 ms [34]. Specifically,  $\pi$ -Orchestrator may suggest different solutions with corresponding slice configurations that translate into variable slice scheduling opportunities per decision interval. Therefore, we evaluate the Cumulative Distribution Functions (CDFs) of traffic latency considering different scheduling opportunities. Despite the resource deficit introduced by the ENS, the average slice latency improves when the number of scheduling opportunities increases. Fig. 9b depicts the temporal evolution of the overall effects caused by the ENS slice setup at transmission time interval (TTI) level. All traffic traces, MCS and channel quality values are generated with a millisecond granularity according to their temporal evolution and statistical distributions resulting from our monitoring dataset. The initial scenario accounts for a fixed resource allocation scheme (without any ENS running) which allows to satisfy all the different slice requirements, as expected in normal traffic conditions thanks to the adoption of calibrated admission and control algorithms, e.g. [28]. In particular, among the 4 slices populating the system, eMBB slices have a fixed quota of 30% of available resources while URLLCs have 20%, respectively. After 10 seconds, an ENS is instantiated. Its higher priority implies the execution of the  $\pi$ -Orchestrator to minimize service disruption. From the upper plot, it can be noticed how our solution reduces the quota of resources



(a) Resource orchestration at system level.



(b) Resource orchestration at base station level.

Figure 9: Performance evaluation of a 4-slices scenario when an ENS is installed.

assigned to non-delay sensitive services which, unavoidably, suffer from a resource deficit (highlighted in red). However, in emergency scenarios the RAN slice scheduler may assign unused reserved resources to active slices with traffic pending to be served, further reducing the resource deficit and, in turn, the SLA violation of certain slices. In the lower plot, we depict the overall traffic potentially violating the SLAs with a red dashed line, whereas the traffic actually exceeding it with a black continuous line. Results show how scheduling relaxation would help in reducing the SLA violations up to 30%. Clearly, this result assumes that none of the resources allocated for ENSs can be consumed by other slices, even in absence of emergency traffic.

## VI. RELATED WORK

The ever-increasing vehicular traffic fosters the need to deeply understand the complex relationship that regulates mobility patterns which, in turn, affects mobile network operational conditions [35], [36]. From the one side, this effort requires constant monitoring of the communication infrastructure. From the other side, the highly heterogeneous set of monitoring KPIs demands for advanced solutions to automate the early detection of anomalous conditions. The authors of [37] initially addressed this scenario proposing a Bayesian network working on a set of discrete metrics collected from an operational UMTS infrastructure. Similarly, the work of [38] focused on a large-scale cellular network scenario, where traffic traces are modeled into regular and random components. Their decomposition approach well suits predictable patterns, but fails in highly variable scenarios. More recently, state-of-the-art solutions start combining monitoring traces and heterogeneous contextual information to improve the effectiveness of model predictions and decisions. In [27], [29] the authors leverage spatio-temporal characteristics of mobile traces to predict resource utilization in the context of network slicing.

Most of the works in the literature address the urban environment and aims to mitigate the cause of traffic con-

gestions by predicting traffic flows and offloading strategies to alternative paths [39], while less studies targeted the highway scenario. In [40] the authors present a spatio-temporal analysis of highways travel patterns exploiting per-vehicle data records collected by a centralized authority in China. Given the fine granularity of per-vehicle information, there is lack of a discussion about anomaly detection from a mobile network perspective. Additionally, note that the majority of the related works presented above tackle the vehicular traffic prediction by means of active systems, e.g., GPS location exchange, or networks of sensors deployed along the road, which clearly ease the prediction task but also increases the volume of information generated along the process. To the best of our knowledge, this is the first work considering the prediction of road events and their impact on the mobile network infrastructure exclusively accounting for *aggregated* and *passive* RAN monitoring samples.

## VII. CONCLUSIONS

In this paper, we presented  $\pi$ -ROAD, a deep learning-based solution for the analysis and detection of road events over the highway. The model accounts for two complementary stages *i*) an autoencoder-based stage to identify anomalous patterns in the temporal evolution of operational mobile network data, and *ii*) a 3D-CNN-based stage to overcome simple threshold detection schemes and automatically learn the relationship that links multiple metric-specific anomalies to road event occurrences. The output of  $\pi$ -ROAD may be used in a multitude of emergency scenarios. We focused on the design and validation of a slice orchestration solution dealing with the setup of Emergency Network Slice (ENS) in V2X scenarios. Considering real mobile traffic distributions from a major highway in Italy, our results show that the information provided by  $\pi$ -ROAD can significantly reduce, up to 30%, the impact of Emergency Network Slice setup, also decreasing the probability of service disruption on other running network slices.

## VIII. ACKNOWLEDGMENTS

The work of Politecnico di Milano was supported by the H2020 ITN Project Spotlight (Grant Agreement No. 722788). The work of NEC Laboratories Europe was supported by the H2020 Mon5G Project (Grant Agreement No. 871780) and by the H2020 Carmen Project (Grant Agreement No. 825012).

## REFERENCES

- [1] Third Generation Partnership Project (3GPP), “Technical Specification Group Services and System Aspects; System architecture for the 5G System (Rel. 16),” 3GPP TS 23.501 V16.5.1, Aug. 2020.
- [2] P. Rost *et al.*, “Network Slicing to Enable Scalability and Flexibility in 5G Mobile Networks,” *IEEE Communications Magazine*, vol. 55, no. 5, pp. 72–79, May 2017.
- [3] J. Ordóñez-Lucena, P. Ameigeiras, D. Lopez, J. J. Ramos-Munoz, J. Llorca, and J. Folgueira, “Network Slicing for 5G with SDN/NFV: Concepts, Architectures, and Challenges,” *IEEE Communications Magazine*, vol. 55, no. 5, pp. 80–87, May 2017.
- [4] W. H. Organization, Jul. 2020. [Online]. Available: {[https://www.who.int/gho/road\\_safety/mortality/en/](https://www.who.int/gho/road_safety/mortality/en/)}
- [5] J. W. K. Donoughe and H. Gabler, “Analysis of Firetruck Crashes and Associated Firefighter Injuries in the United States,” vol. 56, no. 1. Association for the Advancement of Automotive Medicine, Oct. 2012, pp. 69–76, Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3503424/>.
- [6] J. Wang, J. Tang, Z. Xu, Y. Wang, G. Xue, X. Zhang, and D. Yang, “Spatiotemporal Modeling and Prediction in Cellular Networks: A Big Data Enabled Deep Learning Approach,” in *IEEE Conference on Computer Communications - INFOCOM’17*, May 2017.
- [7] C. Marquez, M. Gramaglia, M. Fiore, A. Banchs, and X. Costa-Pérez, “How Should I Slice My Network?: A Multi-Service Empirical Evaluation of Resource Sharing Efficiency,” in *ACM International Conference on Mobile Computing and Networking*, Oct. 2018, pp. 191–206.
- [8] G. Khodabandelou, V. Gauthier, M. Fiore, and M. A. El Yacoubi, “Estimation of Static and Dynamic Urban Populations with Mobile Network Metadata,” *IEEE Transactions on Mobile Computing*, vol. 18, no. 9, pp. 2034–2047, Sept. 2019.
- [9] A. Furno, M. Fiore, R. Stanica, C. Ziemlicki, and Z. Smoreda, “A Tale of Ten Cities: Characterizing Signatures of Mobile Traffic in Urban Areas,” *IEEE Transactions on Mobile Computing*, vol. 16, no. 10, pp. 2682–2696, Oct. 2017.
- [10] Third Generation Partnership Project (3GPP), “Technical Specification Group Services and System Aspects; Policy and charging control architecture (Rel. 16),” 3GPP TS 22.203 V16.2.0, Dec. 2019.
- [11] C. E. Andrade, S. D. Byers, V. Gopalakrishnan, E. Halepovic, D. J. Poole, L. K. Tran, and C. T. Volinsky, “Connected Cars in Cellular Network: A Measurement Study,” in *ACM Internet Measurement Conference - IMC’17*, Nov. 2017, p. 235–241.
- [12] X. Wang, Z. Zhou, F. Xiao, K. Xing, Z. Yang, Y. Liu, and C. Peng, “Spatio-Temporal Analysis and Prediction of Cellular Traffic in Metropolis,” *IEEE Transactions on Mobile Computing*, vol. 18, no. 9, pp. 2190–2202, Sept. 2019.
- [13] B. Liao, J. Zhang, C. Wu, D. McIlwraith, T. Chen, S. Yang, Y. Guo, and F. Wu, “Deep Sequence Learning with Auxiliary Information for Traffic Prediction,” in *ACM International Conference on Knowledge Discovery and Data Mining*, Jun. 2018, pp. 537–546.
- [14] S. Tariq, S. Lee, Y. Shin, M. S. Lee, O. Jung, D. Chung, and S. S. Woo, “Detecting Anomalies in Space Using Multivariate Convolutional LSTM with Mixtures of Probabilistic PCA,” in *ACM International Conference on Knowledge Discovery and Data Mining*, Jul. 2019, pp. 2123–2133.
- [15] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [16] H. D. Trinh, E. Zeydan, L. Giupponi, and P. Dini, “Detecting Mobile Traffic Anomalies Through Physical Control Channel Fingerprinting: A Deep Semi-Supervised Approach,” *IEEE Access*, vol. 7, pp. 152 187–152 201, Oct. 2019.
- [17] Y. Zhu, H. Li, Y. Liao, B. Wang, Z. Guan, H. Liu, and D. Cai, “What to Do Next: Modeling User Behaviors by Time-LSTM,” in *International Joint Conference on Artificial Intelligence*, vol. 17, 2017, pp. 3602–3608.
- [18] Y. Yang, Z. Hu, K. Bian, and L. Song, “ImgSensingNet: UAV vision guided aerial-ground air quality sensing system,” in *IEEE Conference on Computer Communications - INFOCOM’19*, Jul. 2019, pp. 1207–1215.
- [18] S. Ji, W. Xu, M. Yang, and K. Yu, “3D Convolutional Neural Networks for Human Action Recognition,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 1, pp. 221–231, Jan. 2013.
- [20] A. Zappone *et al.*, “Wireless Networks Design in the Era of Deep Learning: Model-based, AI-based, or both?” *IEEE Transactions on Communications*, vol. 67, no. 10, pp. 7331–7376, Jun. 2019.
- [21] K. Sankhe, M. Belgiovine, F. Zhou, S. Riyaz, S. Ioannidis, and K. Chowdhury, “ORACLE: Optimized radio classification through convolutional neural networks,” in *IEEE Conference on Computer Communications - INFOCOM’19*, Jul. 2019, pp. 370–378.
- [22] J. Pérez-Romero *et al.*, “Supporting Mission Critical Services through Radio Access Network Slicing,” in *International Conference on Information and Communication Technologies for Disaster Management*, Dec. 2019, pp. 1–8.
- [23] L. Zanzi, V. Sciancalepore, A. Garcia-Saavedra, X. Costa-Pérez, G. Agapiou, and H. D. Schotten, “ARENA: A Data-driven Radio Access Networks Analysis of Football Events,” *IEEE Transactions on Network and Service Management*, pp. 1–14, Oct. 2020.
- [24] Third Generation Partnership Project (3GPP), “Study on enhancement of 3GPP Support for 5G V2X Services (Release 16),” 3GPP TR 22.886 V14.0.0, Dec. 2018.
- [25] —, “Feasibility Study on New Services and Markets Technology Enablers, Stage 1 (Release 14),” 3GPP TR 22.891 V14.2.0, Sep. 2016.
- [26] J. Mei, X. Wang, and K. Zheng, “Intelligent Network Slicing for V2X Services Toward 5G,” *IEEE Network*, vol. 33, no. 6, pp. 196–204, 2019.
- [27] D. Bega, M. Gramaglia, M. Fiore, A. Banchs, and X. Costa-Pérez, “DeepCog: Cognitive Network Management in Sliced 5G Networks with Deep Learning,” in *IEEE Conference on Computer Communications - INFOCOM’19*, Apr. 2019, pp. 280–288.
- [28] J. X. Salvat, L. Zanzi, A. Garcia-Saavedra, V. Sciancalepore, and X. Costa-Pérez, “Overbooking Network Slices Through Yield-driven End-to-end Orchestration,” in *ACM CoNEXT*, Nov. 2018, pp. 353–365.
- [29] D. Bega, M. Gramaglia, M. Fiore, A. Banchs, and X. Costa-Pérez, “AZTEC: Anticipatory Capacity Allocation for Zero-Touch Network Slicing,” in *IEEE Conference on Computer Communications - INFOCOM’20*, Jul. 2020, pp. 280–288.
- [30] Weihua Wang, P. B. Luh, J. H. Yan, and G. A. Stern, “An Improved Lagrangian Relaxation Method for Discrete Optimization Applications,” in *IEEE International Conference on Automation Science and Engineering*, Aug. 2008, pp. 359–364.
- [31] S. Martello and P. Toth, *Knapsack Problems: Algorithms and Computer Implementations*. USA: John Wiley Sons, Inc., 1990.
- [32] J. Foulkas, M. K. Marina, and K. Kontovasilis, “Orion: RAN Slicing for a Flexible and Cost-Effective Multi-Service Mobile Network Architecture,” in *ACM International Conference on Mobile Computing and Networking*, Oct. 2017.
- [33] L. Zanzi, V. Sciancalepore, A. Garcia-Saavedra, H. D. Schotten, and X. Costa-Pérez, “LACO: A Latency-Driven Network Slicing Orchestration in Beyond-5G Networks,” *IEEE Transactions on Wireless Communications*, pp. 1–15, Oct. 2020.
- [34] X. Ge, “Ultra-Reliable Low-Latency Communications in Autonomous Vehicular Networks,” *IEEE Transactions on Vehicular Technology*, vol. 68, no. 5, pp. 5005–5016, Mar. 2019.
- [35] R. Li, Z. Zhao, J. Zheng, C. Mei, Y. Cai, and H. Zhang, “The Learning and Prediction of Application-Level Traffic Data in Cellular Networks,” *IEEE Transactions on Wireless Communications*, vol. 16, no. 6, pp. 3899–3912, Jun. 2017.
- [36] J. Zhang, Y. Zheng, J. Sun, and D. Qi, “Flow Prediction in Spatio-Temporal Networks Based on Multitask Deep Learning,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 32, no. 3, pp. 468–478, Mar. 2020.
- [37] R. M. Khanafer, B. Solana, J. Triola, R. Barco, L. Moltsen, Z. Altman, and P. Lázaro, “Automated Diagnosis for UMTS Networks Using Bayesian Network Approach,” *IEEE/ACM Transactions on Vehicular Technology*, vol. 57, no. 4, pp. 2451–3058, Oct. 2008.
- [38] F. Xu, Y. Lin, J. Huang, D. Wu, H. Shi, J. Song, and Y. Li, “Big Data Driven Mobile Traffic Understanding and Forecasting: A Time Series Approach,” *IEEE Transactions on Services Computing*, vol. 9, no. 5, pp. 796–805, Aug. 2016.
- [39] C. Ide *et al.*, “LTE Connectivity and Vehicular Traffic Prediction Based on Machine Learning Approaches,” in *IEEE Vehicular Technology Conference (VTC-Fall)*, Sept. 2015, pp. 1–5.
- [40] W. Ji, Z. Lu, and T. Zhu, “Mining Spatial-Temporal Travel Patterns from Highway Transaction Data,” in *ACM Symposium on Computer Science and Intelligent Control*, Sept. 2018.