

Measuring the Confidence of Single-Point Traffic Forecasting Models: Techniques, Experimental Comparison, and Guidelines Toward Their Actionability

Ibai Laña¹, Member, IEEE, Ignacio Olabarrieta², and Javier Del Ser, Senior Member, IEEE

Abstract—The estimation of the amount of uncertainty featured by predictive machine learning models has acquired a great momentum in recent years. Uncertainty estimation provides the user with augmented information about the model’s confidence in its predicted outcome. Despite the inherent utility of this information for the trustworthiness of the user, there is a thin consensus around the different types of uncertainty that one can gauge in machine learning models and the suitability of different techniques that can be used to quantify the uncertainty of a specific model. This subject is mostly non-existent within the traffic modeling domain, even though the measurement of the confidence associated to traffic forecasts can favor significantly their actionability in practical traffic management systems. This work aims to cover this lack of research by reviewing different techniques and metrics of uncertainty available in the literature, and by critically discussing how confidence levels computed for traffic forecasting models can be helpful for researchers and practitioners working in this research area. To shed light with empirical evidence, this critical discussion is further informed by experimental results produced by different uncertainty estimation techniques over real traffic data collected in Madrid (Spain), rendering a general overview of the benefits and caveats of every technique, how they can be compared to each other, and how the measured uncertainty decreases depending on the amount, quality and diversity of data used to produce the forecasts.

Index Terms—Uncertainty estimation, confidence, traffic forecasting.

I. INTRODUCTION

ROAD traffic forecasting has been a subject of intense academic research for decades [1]. This modeling task was originally approached based on regressive techniques for time series analysis [2], [3], [4], veering over time towards Machine Learning (ML) methods [5]. When compared to

traditional regression techniques, ML allows for the detection and characterization of hidden relationships among data in exchange for a higher level of complexity. A glance at the recent literature contributed on ML-based traffic forecasting reveals that deep neural networks (also referred to as *Deep Learning*) conform nowadays the core of almost any traffic forecasting work [6], [7], yielding a torrent of works that are not only abundant, but increasing steadily every year [8], [9]. This sustained interest of the research community in this topic suggests that traffic forecasting still poses challenges for the method used to address this modeling task. However, several comprehensive studies expose that traffic forecasting has reached its modeling performance asymptote, wherein the addition of extra layers of complexity in the modeling proposal yields negligible performance gains [10], [11]. This prevailing performance-driven pursuit for new traffic forecasting models is far from being actionable from the perspective of Intelligent Transportation Systems [12]. As a result, most of the challenges of this research area identified years ago remain unsolved to date [1], [13].

This noted lack of actionability in current traffic forecasting solutions can be addressed by moving the research focus away from performance metrics towards augmenting their output with valuable information to support practical decision making processes of traffic managers. When the quality of estimated forecasts reaches enough quality in terms of model predictive error, it is of utmost practical interest to explore the individual level of *confidence* associated to each forecast. Model error is a quantitative measure of the performance of the model when producing forecasts for a certain set of input traffic values [14], [15]. However, both the input data and the modeling technique itself may be subject to different sources of uncertainty [16]. This means that even models with high performance metrics can yield predictions that are inaccurate beyond the threshold within which they are useful [17], which propagates to a confidence level of the model in its predicted outcomes. It is not possible to gauge the error of a given prediction until the actual value (*ground truth*) occurs. However, measuring the uncertainty under which the prediction is furnished – in other words, assigning a confidence level to the prediction – can be approached at inference time using different strategies. Thus, a manager or service consuming such forecasts would

Manuscript received 20 April 2023; revised 24 November 2023 and 18 January 2024; accepted 7 March 2024. Date of publication 18 March 2024; date of current version 29 August 2024. This work was supported in part by European Union’s Horizon 2020 Program through the project “URBANITE: Supporting the decision-making in urban transformation with the use of disruptive technologies” under Grant 870338 and in part by Basque Government (The Consolidated Research Group MATHMODE) under Grant IT1256-22. The Associate Editor for this article was X. Ma. (*Corresponding author: Javier Del Ser*.)

Ibai Laña and Ignacio Olabarrieta are with TECNALIA, Basque Research and Technology Alliance (BRTA), Derio, 48160 Bizkaia, Spain.

Javier Del Ser is with the Department of Communications Engineering, University of the Basque Country (UPV/EHU), 48013 Bilbao, Spain (e-mail: javier.delser@ehu.eus).

Digital Object Identifier 10.1109/TITS.2024.3375936

not only be informed about the estimated traffic value for a certain location in space and time, but also the confidence under which the predicted value can be trusted. Besides, from the practitioners point of view, knowing the uncertainty also provides information about when to stop trying to improve a model.

In this regard, many disciplines (especially those where decisions made based on the model's output may entail a risk for the human life) have dealt with the quantification of uncertainty in data-based modeling, so that the provision of *augmented predictions* with confidence estimates has already become a matter of intense research. Within the medical domain, for instance, the confidence of ML models is crucial in model-based clinical diagnosis [18], and it has been obtained for diverse purposes [19], [20], [21], [22]. Other fields for which the actionability of forecasts is highly relevant have also considered measuring the confidence of predictions, e.g. in the energy sector [23], [24] or in weather and climate modeling [25], [26], [27].

The ITS domain is not unfamiliar to the application of confidence estimation techniques, specially in cases where predictions issued by a model play an essential role in decision making processes. For instance, within the autonomous driving field, vehicles need to sense and anticipate the contextual circumstances in which they operate to determine their behavior in uncertain environment. Researchers working on vehicular perception are highly focused on dealing with uncertainty when estimating the next maneuver [28], [29], [30] or the trajectory of other vehicles and pedestrians [31], [32], [33]. It is also the case of air traffic management, a field comprising multiple subareas in which decisions must be made fast, and errors can entail important operational consequences and high economical costs. As a result, assessing the uncertainty of forecasts has been addressed in this area for diverse purposes, from aircraft trajectory prediction [34] to the demand estimation of airport facilities [35], [36] to help managers scale services. In relation to the latter example, and closer to the road traffic domain, a considerable body of literature related to forecasting uncertainty has focused on travel demand estimation. For instance, authors in [37], [38], [39] examine the statistical robustness of forecasts and propose different methods to detect them. In [40], infrastructure capacity constraints are considered to propose a new methodology that quantifies uncertainty, similar to what [41] does for an established and operating travel demand system in Sweden. Analogous research is performed for transport mode choice predictions in [42], with equivalent actionable results.

All these research contributions deal with the prediction of travel demand, which is intended to be provided to public authorities, infrastructure and public transportation managers and investors. By informing them about the uncertainty associated to the demand estimation, insights on how traffic and the demand of transportation services will operate in the long term are produced, which can be used to scale appropriately such infrastructures and services. When decision making is based on short-term traffic forecasts, the impact of a forecasting

error has a shallower relevance. Anyhow, if these forecasts are meant to be used for any real-world purpose, measuring their uncertainty is of paramount importance for their trustworthiness and actionability. Following up the work in [43], the performance of travel demand forecasts was not properly considered in terms of the uncertainty it was subject to. Up to this point, all the mentioned works involve forecasting up to several minutes or hours. Long-term traffic estimations can also be studied for traffic demand analysis, goal for which measuring their uncertainty becomes an essential actionability driver [44].

This manuscript finds its motivation in the lack of an unified referential work consolidating the state of the art related to uncertainty estimation applied to short-term traffic forecasting problems. To cover this niche, we herein offer a short yet instructive summary about the different types of uncertainty that may arise in road traffic forecasting scenarios, an enumeration of the practical reasons why the confidence of traffic forecasting models is a key driver for their practicality, different uncertainty estimation approaches used nowadays to quantify such a confidence, and scores utilized to compare among such techniques. This multi-faceted analysis is complemented by an extensive experimental benchmark with real traffic and weather data collected over the city of Madrid (Spain). Specifically, the experimental setup and the results obtained therefrom permit to answer with empirical evidence three research questions (RQs) that touch the core of the overall study:

- *RQ1 (Techniques and comparison framework)*: How do different uncertainty quantification techniques perform when applied to a particular traffic forecasting model? Which are their main strengths and weaknesses for each problem? Which scores can be used to compare these techniques? What does each of them contribute to the improvement of the model's actionability?
- *RQ2 (Scenario under study)*: Which changes in the available data affect uncertainty? How can the confidence of a forecasting model be improved by changing/augmenting data at its input?
- *RQ3 (The relevance of calibration)*: What impact does the calibration process of some of the uncertainty estimation techniques have on their outcome? Why is it relevant for traffic forecasting?

The rest of this paper is structured as follows: first, Section II introduces the readership to the different sources of uncertainty that one can encounter in traffic forecasting scenarios. This section also enumerates different practical purposes for which the uncertainty of traffic forecasts must be measured, and provides an overview of uncertainty estimation techniques, along with scores that can be used to compare their effectiveness. Section III introduces the data and the experimental setup devised to answer the RQs formulated above. Section IV presents and discusses the obtained results, structuring their analysis in terms of the RQ under target. Lastly, insightful concluding remarks are offered in Section V, together with future research directions departing from this study.

II. ESTIMATING THE CONFIDENCE OF TRAFFIC FORECASTING MODELS: SOURCES OF UNCERTAINTY, PRACTICAL PURPOSES, TECHNIQUES AND MEASURES

As argued in the introduction, uncertainty quantification can improve significantly the worth of any forecasting model. Although it is an old area of study in statistics and probability, its application to machine learning techniques is relatively recent [45], [46]. It has been exposed before that decision makers in many application fields (e.g., medical diagnosis, industrial prognosis) demand the provision of quantitative metrics related to the trustworthiness and reliability of predictions issued by machine learning models. This is certainly not the case of traffic forecasting, a vibrant field with hundreds of publications on a yearly basis, in which confidence measurements are rarely provided. There are some notable exceptions that were surveyed in [16], which revealed that most works at the time reporting confidence intervals obtained them experimentally by repeating experiments and measuring the variance and standard deviation of forecasts. Some other works published thereafter took advantage of particular characteristics of Kalman filter modeling techniques to extract confidence intervals for the predictions [47], [48]. However, this is not a common practice: in the traffic forecasting literature, only performance scores are usually computed.

Interestingly, machine learning techniques for predictive modeling have reached ever-growing levels of complexity over the years, rendering their modeled knowledge more opaque and difficult to understand for the user consuming their predictions [10]. As a result, different methods for explaining their inner structure in an user-interpretable manner are under active investigation lately [49]. As a supplement to such explainability methods, informing about the uncertainty associated to the predictions supports even further the trustworthiness of forecasting models, and favors decision making processes not only linked to the usefulness of the traffic forecasts themselves (e.g. traffic congestion management), but also in regards to the intelligent collection of traffic data.

This section delves into these ideas by introducing briefly to the different sources of uncertainty in traffic forecasting (Subsection II-A), by discussing on the diverse purposes and inherent usefulness of uncertainty estimation in this context (Subsection II-B), by shortly reviewing state-of-the-art uncertainty estimation techniques (Subsection II-C), and by enumerating different quantitative metrics that can be used to evaluate the quality of estimated confidence intervals (Subsection II-D). The overarching goal of this section is to set the knowledge basis towards the experimental part of this work.

A. Sources and Types of Uncertainty in Traffic Forecasting

The uncertainty of a model trained with supervised learning can have different kinds of sources. It is commonly accepted that this uncertainty of the output of the model (evaluations of the model with new data) can include a fraction of *epistemic uncertainty*, that accounts for all the uncertainty that the modeling process itself introduces, and *aleatory uncertainty*, that regards the irreducible uncertainty that is present in the




Fig. 1. Sources of uncertainty present in traffic forecasting.

data due to their inherent variability [50]. However, both the number of different uncertainty sources and their *separability* are open practical questions subject to debate among statisticians, frequently an outlet for deeper philosophical discussion [51]. Epistemic uncertainty can (theoretically) be reduced, being it part of the way in which data are structured and represented. On the contrary, aleatory uncertainty cannot be decreased, as it is inherent to the variability and noisy nature of data. Discerning among both sources of uncertainty can be interesting precisely for reducing the share of epistemic uncertainty. Nonetheless, burrowing through the philosophical grounds of the very definition of uncertainty [52], it is unclear whether different sources of uncertainty can be disentangled from each other, or even that the uncertainty associated to data can be isolated from other data features.

Practical implications of segregating aleatory and epistemic uncertainty in traffic forecasting are clear: while epistemic uncertainty can be reduced to an extent depending on the modeling choice in use, measuring the amount of aleatory uncertainty can dictate whether the addition of new variables and/or the collection of new data helps reducing it effectively. Unfortunately, from a practical point of view such a segregation of uncertainty types might not be feasible [51]. Moreover, it might be even irrelevant as long as mechanisms are designed to reduce the *global* uncertainty that aggregates both sources. In fact, not all uncertainty estimation techniques discussed in this work provide separated uncertainty measurements, consideration that will be contemplated in their experimental comparison.

B. Uncertainty Estimation in Traffic Forecasting: What For?

The use of uncertainty estimation techniques in some of the aforementioned fields is straightforward: quantifying the confidence under which a model produces its output is critical for its actionability in many diverse applications. For this reason, research about uncertainty estimation methods is growing and becoming available for very diverse machine learning models. Short-term traffic forecasting has probably a less critical nature than, for instance, medical or finance sectors. However, measuring uncertainty of traffic forecasts can also be profitable if we consider the purposes for which it can be performed. Such purposes are graphically illustrated in Figure 2 and explained in what follows:

- *Improved actionability*: the unequivocal application of estimating uncertainty is to provide practitioners with a

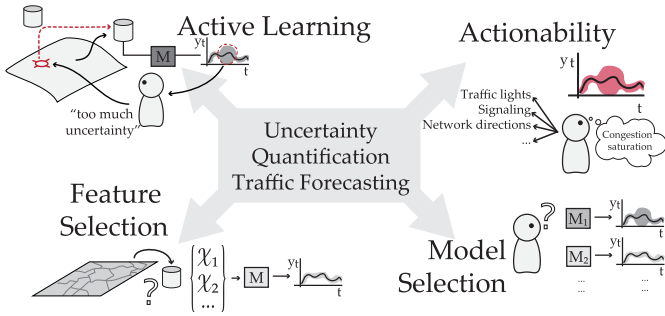


Fig. 2. Purposes of uncertainty estimation in the context of traffic forecasting.

confidence interval of predictions, so that better decisions can be made [12]. Unlike customary regression scores, which are commonly obtained for a whole test dataset and apprise about the general (averaged) performance of the trained model, uncertainty metrics reach the individual prediction level, giving confidence information for each data point predicted by the model. This not only provides valuable information about the extent to which each individual prediction can be trusted, but also allows for a further study on how traffic behaves in different time intervals and seasons over the year.

- Model selection and comparison:** as explained before, a part of the total uncertainty of the output of a forecasting scheme can be attributed to the way the information is modeled. Uncertainty can be, therefore, another dimension in a model comparison framework, allowing researchers and ITS experts to choose models that perform with narrower confidence intervals. Models performing similarly in terms of regression error but introducing less epistemic uncertainty, can be considered more desirable in some contexts. Furthermore, examining the uncertainty associated to individual predictions can be also interesting to select models that reduce the uncertainty in critical segments of the traffic data series (e.g. in hours featuring high traffic intensity).
- Feature selection:** prior to the application of a training algorithm that fits the parameters of the model to represent a reality represented by data, raw traffic information is arranged into a dataset with a feature space \mathcal{X} and an output space \mathcal{Y} . Even before the training process starts, the way in which variables are selected and processed for constructing \mathcal{X} and \mathcal{Y} can affect their degree of uncertainty [51]. This uncertainty should be deemed as epistemic, as it falls under the way knowledge is represented and usually is subject to the experience of the expert at hand. From this conceptual point of view, it is interesting to highlight that the way in which input data are selected, shaped and fed to the model towards its training must be also considered a *part of the model*. Thus, analogously to the previous concept, uncertainty metrics can help, within the framework of a benchmark, to define which input features provide more trustworthy results. Likewise, metrics can be helpful to seek features that help reducing the uncertainty induced by the overall forecasting model.

- Active learning:** uncertainty estimation is also an essential component of algorithms that learn actively [53]. Active learning consists of constructing training datasets dynamically, where training samples are added progressively and aimed at reducing uncertainty and increasing diversity. Latest research in this vibrant field is highly reliant on uncertainty estimation [54], [55], and the way in which uncertainty is measured is a matter of study in itself [56]. In the context of road traffic forecasting, the installation of new traffic reading sensors (either provisional or fixed) can be decided upon the study of the distribution of uncertainty in space and time, especially when traffic measurements are scarce in locations of interest for traffic management. If forecasts for such locations exhibit a higher degree of uncertainty, new measurement campaigns can be commanded to collect more data that allow improving the confidence of the predictions issued by the model. This is especially relevant in practical circumstances where traffic profiles for a certain place are inferred from data collected in other nearby points of the road traffic network [57].
- Trend change detection:** a well-calibrated confidence interval for the output of the model should characterize, with the established confidence level, the percentage of real traffic data that falls within the interval. This *a priori* notion can be used to detect whether, in a traffic data stream, an excessive number of real traffic data points fall outside the confidence interval of the model. If this occurs, it can be an evidence of a change in the underlying distribution of traffic data requiring, upon its detection, an update of the model's knowledge (via incremental training, selective forgetting or any other adaptation method alike). Despite its unquestionable practical benefits (especially in dynamic urban road networks prone to different non-stationarities affecting its traffic flows), this adaptation approach have not been studied until very recently [58].

There are deeper questions related to uncertainty that could concern practitioners, such as the risk of introducing dependence among random events that affect the system being modeled [51]. Nonetheless, the essential actionability aspects presented above can help ITS stakeholders design more confident models, and understand them better in order to put their output to practice.

C. A Brief Overview of Uncertainty Estimation Techniques

Quantifying the uncertainty present in the output of a model implies measuring the way in which it can vary. This can be expressed as the variance or standard deviation of each predicted value, or by reporting different statistics (e.g. confidence intervals or percentiles) of the estimated output distribution [16]. When dealing with machine learning models, once a dataset is built and the parameters of the chosen model are trained through a training algorithm, the outputs of such model will not vary given a fixed input. For a certain forecasting query, there is only one output given one input to the model. Unless probabilistic formulations of the model's

parameters are formulated (as in Bayesian neural networks), most machine learning models are *deterministic* after their parameters have been trained. Thus, with an already trained model, estimating the variability of each forecast requires in general an uncertainty estimation technique.

On this basis, many techniques and methods have been developed by the community to characterize the output distribution of a machine learning model based on prior knowledge from training data. Once this output distribution is estimated, the usual approach consists of defining significance levels and obtain confidence intervals. Some machine learning models allow changing their default operation to obtain predictions based on percentiles. These model-specific approaches have been described in the literature as *intrinsic* [59]. Conversely, *extrinsic* uncertainty estimation techniques are not specific to a particular machine learning model, but allow obtaining a measurement of uncertainty once a model has been trained (i.e., a post-hoc estimation), normally through a calibration process.

Even though some authors have tried to categorize the landscape of uncertainty measuring techniques [59], it appears to be difficult to find consensus beyond the differences between intrinsic and extrinsic techniques exposed above. For instance, some techniques allege to be able to discriminate between the epistemic and aleatory uncertainty present in the output of machine learning models [60], [61], frequently based on Bayesian formulations of their underlying training mechanisms. Other branch of techniques based on wrapper methods rely on a calibration process prior to the training phase of the algorithm, whereas other techniques hinge on the use of ensembles or the generation and modeling of multiple data bootstraps to approximate the output distribution. Additionally, different variants of deep neural networks have been proposed in the last years to estimate the output distribution, via the probabilistic definition of their trainable weights, evidential formulations of the loss function used to guide the training process, or the derivation of mechanisms to sample the sought distribution. The more recent are the research works in this, the more specific for deep learning they happen to be. In [62], an uncertainty estimation framework is proposed, which focuses only on deep spatio-temporal forecasting. Methods are very similar to the ones proposed in other works: methods that produce distributions in their inner behavior (Bayesian approaches), and methods which are modified so the probability distribution of their output is forced (frequentist approaches). A similar characterization is proposed in [63].

Beyond the discussion on the suitability of one taxonomic criterion or another, in what follows we examine the most widely adopted options to estimate the uncertainty in machine learning models used for regression tasks. To this end we pause briefly at the specific techniques organized in the taxonomy shown in Figure 3. This short review connects tightly with the experiments designed to answer RQ1, where a comprehensive comparison benchmark over real traffic data will expose the strengths, weaknesses and the applicability of the methods reviewed below:

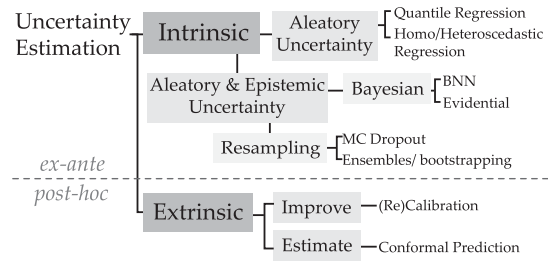


Fig. 3. Taxonomy of uncertainty estimation techniques tackled in this work.

- *Conformal prediction*: as defined in [64], [65], conformal predictors are *confidence* predictors. Their operation is based on the conception of *nonconformity* measures, which allow predicting regions of points instead of individual points (i.e., confidence intervals) by using the statistical knowledge that can be obtained through a calibration process from the training samples. This method has been prevalent in the machine learning community for years [66], being used in applications arising in many fields [67]. Conformal prediction is agnostic to the particularities of the underlying model, and can be used, in principle, with any machine learning algorithm and dataset configuration.
- *Ensemble methods*: ensemble learning consists of the aggregation of knowledge obtained from different models learned from different *views* of the data. Combining the predictions issued by different methods learned from the same data has been proven to yield better performance than any contributing member to the ensemble, which is achieved by reducing the spread in the predictions made by such models [68], [69]. Besides the increased robustness favored by this smoothing process, ensembles represent an easy and straightforward approach to sample the distribution associated to the combined output; once trained, each base learner yields a different prediction for a given input, all of which are aggregated by averaging or other more sophisticated combination procedures (e.g. weighting based on out-of-bag performance estimates). Thus, confidence intervals can be obtained by characterizing statistically the distribution of all the different predictions produced by each learner in the ensemble for a given input [70].
- *Bootstrapping*: this model-agnostic technique is used to estimate the uncertainty associated with a sample statistic by resampling with replacement from the original dataset (*bootstrapping*). All such bootstraps are used to learn a forecasting model. The outputs of such forecasting models for a given test instance are collected and used to create a sampling distribution of the statistic of interest (e.g., the traffic forecast), which represents the variability or uncertainty associated with the estimate. From this distribution we can compute prediction intervals, i.e. a range within which the true traffic measure is likely to fall [71]. In addition to different works relying on bootstrapping for estimating the uncertainty of different

traffic data models [40], [72], [73], modifications of the naïve bootstrap procedure have been proposed to account for all sources of uncertainty in prediction intervals. One of such methods [74] will be used for the benchmark discussed in the experimental part of this manuscript.

- *Quantile regression*: This family of methods was originally conceived for linear regression methods, for which, instead of using the median as an output, asymmetric errors were considered in order to have upper and lower boundaries of an interval defined by a quantile [75]. More modern contraptions have used this notion as a base to estimate uncertainty of other machine learning models [76], [77] or deep learning models in which the different quantiles are predicted directly with the trained model, as an output of the model [78]. Quantiles have been introduced in loss functions of several learning algorithms, from gradient boosting regression to deep neural networks [79]. These reformulation of the loss functions permit to predict not the median, but instead the extremes of an interval defined by the quantile sought. Their main difference with respect to previous methods in terms of prediction interval is the distribution of points inside the interval. Other methods generate predictions that may have an unknown distribution. Once they are characterized, the boundaries of the interval are obtained according to a significance level. This approach *produces only boundary lines*, and they are normally distributed. A characteristic prediction would be the median, but any other *predicted* points inside the interval would be placed according to their significance level. This becomes a relevant issue when comparing the output of these methods to other uncertainty estimation approaches, as the experimental discussions later held will clearly show.
- *Bayesian inference based algorithms*: in general, these algorithms aim to characterize the posterior probability distribution of the model parameters based on the available data and the assumption of an a priori distribution of the parameters. Once this distribution is characterized after training, we have the possibility of establishing confidence levels. Particularly with deep Bayesian networks, the parameters of the model (weights) are driven by a priori probability distribution (often set to be Gaussian), making the so-called family of Bayesian neural networks a model choice particularly suited to estimate uncertainty in modeling problems [80]. Many authors in different fields have taken advantage of Bayesian formulations. For instance, a thorough review of different applications of Bayesian learning approaches in the health domain can be found in [81], whereas use cases have been reported for other diverse scenarios such as plant disease detection [82], solar power forecasting [83] or weather forecasting [84], among others. More recently, some attempts at applying Deep Bayesian neural networks for the delivery of confidence-aware road traffic forecasts have been done [85], [86], but efforts in this regard to date have been significantly lower than in other application scenarios.
- *Model-specific methods*: finally, in this last category we gather all techniques that are specific to a certain family of forecasting models. Many of them have been proposed for deep neural networks, leveraging the flexibility granted by the use of a loss function to guide the training process, or the existence of neural computation mechanisms that allow for an approximate representation of the sought output distribution. Salinas et al. [87] propose modelling the conditional distribution of the upcoming samples of the time series, for which a likelihood distribution is assumed, and the predictions produced by the trained model are samples of such distribution, analogously to Bayesian methods. A representative technique of the former approach is *evidential deep learning* [88], which was proposed to overcome the high computational complexity of the training process of Bayesian neural networks and the dependence of the estimated uncertainty on the suitability of the assumed priors for the network weights. In doing so, evidential deep learning proposes a framework in which a prior distribution is placed on the statistics of a Gaussian output distribution (as opposed to Bayesian neural networks, where priors are imposed to the network weights), so that the addition of newly observed traffic samples provides more support for the neural network to learn the parameters of the evidential distribution. On the other hand, Montecarlo Dropout [89] resounded loudly in the community for its simplicity and scalability to integrate the model's output likelihood by randomly switching off neurons in a neural network at inference time. In a recent contribution, authors of [62] identified novel methods based on imposing variability of the model output and thus characterizing its distribution. On one hand, Mean Interval Score regression uses the Mean Interval Score (MIS), also known as Winkler score [90] as the loss function to minimize in a neural network learning process. This score is based on the interval length at each point, and its computation is very similar in practice to quantile regression, as the upper and lower intervals of the prediction are predicted individually once a significance level is defined. Thus, a neural network aimed at reducing the distance between the intervals that guarantees the significance levels would return the fittest intervals. On the other hand, and in the line of making small perturbations to a deep neural network in order to get variability in the outcome, authors have also resorted to the model-agnostic bootstrapping approach commented previously [91]. Lastly, an *intrinsic* approach included in IBM's UQ360 library [59] is *heteroscedastic regression*, which takes advantage of the expected noise of the data in the model to capture both kinds of uncertainty, as it was shown in [92]. An alternative is *homoscedastic regression*, which assumes that noise is constant across data points. Additional uncertainty estimation techniques for deep learning models were gathered in a recent review [63]. The principles underlying most techniques reviewed therein remain the same: to infer the output distribution with probabilistic models (different flavours of Bayesian models), or force the model to have one

with a frequentist approach like ensembling, Montecarlo Dropout, *MIS* regression or bootstrapping.

- *Methods suited for graph-based models*: traffic variables often hold spatio-temporal correlations that can be exploited for forecasting when dealing with short prediction horizons, geographically close measurement locations and moderate to high sampling rates. When this is the case, graph neural networks have received much attention in the last few years due to their inherent capability to capture spatial relationships by conceiving the road network as a graph. Their combination with convolutional and/or recurrent neural blocks have been proven to excel at modeling such spatio-temporal correlations [9]. When it comes to the estimation of spatio-temporal forecasts provided by these models, proposals have been scarce and mostly restricted to extensions of the above model-specific methods for Deep Learning models. This includes mostly Bayesian formulations of graph neural networks [93], [94], [95] or ensembles [96], [97], [98], [99]. More recently, evidential deep learning has been explored for spatio-temporal traffic modeling using graph neural networks [100].

D. Measuring and Comparing Uncertainty Estimations

The output of any of the methods presented above should be regarded as a measure of the amount of uncertainty associated to their predictions. This is not a trivial aspect: as argued in Subsection II-A, boundaries between different types of uncertainty are not clear, and not all methods measure the model's confidence by following the same principles. For this reason, there is a series of metrics that, departing from a common ground, not only provide a notion of the amount of uncertainty, but also indicate whether the measurement of uncertainty itself is performed equally. As in any other regression task, in traffic forecasting the measurement of the uncertainty at the output of a model can be performed by establishing a significance level beforehand (usually denoted as α), and by defining confidence intervals that contain a fraction of the possible outputs that matches the defined α . Thus, uncertainty relates to the *width* of the interval for each estimated point. For a certain significance level α , a narrower interval represents a less uncertain output of the model for which it is estimated.

Beyond these baseline principles, there is no consensus in the literature regarding these metrics, to the point of referring to the same metric with different names, or using metrics that cannot be measured for all uncertainty estimation methods. With a practical stance, we now summarize a list of uncertainty metrics that can be used to measure different aspects of the estimated uncertainty, their different usages and practical implications. The adopted nomenclature for the rest of the paper is the one used in [101]:

- *Interval width*: this first metric refers to the amplitude between the lowest and highest points of the estimated interval around each predicted data point. The concept of *efficiency* or *informational efficiency* of a conformal predictor [65] is related to this width, while a common

uncertainty quantification metric known as *sharpness* [102], which is computed based on the variance at each point, measures essentially the same. The name *informational efficiency* used by some sources in the related literature refers to the way in which the interval informs: a very narrow interval with high confidence is very informative, while a very wide interval that covers a great number of possible values gives less information with the same confidence, thus it is less efficient. This metric can be also found referred to as *Mean Interval Length* (MIL, as in [101], [103]), or *Mean Prediction Interval Width* (MPIW, [59], [104]). Given a significance level α , this metric expresses how wide must the interval be to include the fraction of values defined by α . As the width can be different for each traffic forecast, it is common to obtain an averaged value of MIL (or MPIW). Mathematically, this metric is defined as:

$$\text{MIL} = \frac{1}{T} \sum_{t=1}^T (u_t - l_t), \quad (1)$$

where u_t and l_t are the upper and lower boundaries of the confidence interval estimated by the technique at hand, and T is the number of instances in the test dataset.

- *Interval coverage*: this second metric is the fraction of the possible outputs that is covered by the interval. This can be related to the so-called *validity* metric [65], and has been also referred to as *Interval Coverage Percentage* (ICP, [101], [103]) or *Prediction Interval Coverage Probability* (PICP, [59], [104]). This metric provides a notion of how valid the interval is, as its value should concur with the percentage of samples defined by the established significance level $(1 - \alpha)$. The metric is defined as per eq. 2:

$$\text{ICP} = \frac{1}{T} \sum_{t=1}^T \mathbb{I}(l_t \leq y_t \leq u_t), \quad (2)$$

where l_t and u_t are the lower and upper confidence interval boundaries as in Expression (1), y_t denotes the ground truth value associated to the t -th predicted value \hat{y}_t in the test set, and $\mathbb{I}(\cdot)$ denotes an auxiliary binary function taking value 1 if its argument is true (0 otherwise).

- *Interval width with relation to the forecasting error*: RMIL [101] is a variant of MIL that relates the size of the estimated confidence interval to the error of the forecast. This allows for larger intervals in order to cover those cases with largest forecasting error (i.e., those traffic values that result to be more difficult to forecast precisely). As a result, RMIL permits to compare two forecasting models that deal with different input predictors in terms of uncertainty. The metric is defined as:

$$\text{RMIL} = \frac{1}{T} \sum_{t=1}^T \frac{(u_t - l_t)}{|y_t - \hat{y}_t|}, \quad (3)$$

where y_t is the real value for the t -th test instance, and \hat{y}_t denotes the predicted value for that instance issued by the forecasting model. This metric can return very high

(even infinite) values if the prediction is very close to the real values, reason for which it has not been considered for the experimentation of this work.

- *Calibration curves*: The calibration of a forecasting model aims to achieve a statistical consistency between the distribution of the forecasts and the distribution of the real values [105]. The outputs of a well-calibrated forecasting model will follow a similar distribution to the real values of the ground truth the model is representing. This property is explored through calibration curves, i.e., a two-dimensional plot relating every predicted value to the real observation it approximates [59]: as such, the optimal calibration curve is the identity function, whereas the deviation of the curve with respect to the optimal calibration is denoted as *calibration error*. It is relevant to note that, although previous metrics can be comparable for any uncertainty estimation approach, the way in which output distributions are obtained may affect substantially to the measurement of the calibration curve and error.
- *Other metrics*: Gneiting et al. [106] proposed a set of *proper scoring rules* that cover, among others, diverse probabilistic aspects of interval forecasts, including combined metrics of some of the aforementioned scores, e.g., *calibration* and *sharpness*. Wu et al. [62] uses one of these scoring metrics, the Mean Interval Score MIS, not only as an evaluation, but as objective function to minimize in order to obtain confidence estimates. Metrics like Negative Log Likelihood, Interval Score or Check have been used in a variety of uncertainty estimation works [107], providing a different perspective to the analysis of the size of the intervals and their coverage of real samples. Considering that they provide similar insights as the above metrics, and there are some hindrances to apply them to all of the uncertainty estimation methods (for instance, some of them are only intended for quantile-based methods), we have prioritized the metrics introduced previously in our experiments for the sake of a clearer and more uniform analysis.

These metrics can be used to assess the impact of different datasets and configurations of the forecasting model, as well as to examine the convenience of each of the uncertainty quantification methods. The insights extracted from their analysis will be helpful answering the different research questions of the experimental study discussed in the next section.

III. EXPERIMENTAL SETUP

In order to assess and compare the uncertainty estimation techniques and confidence metrics presented above, and to discuss on their implications for the actionability of traffic forecasts, an experimental setup has been defined, which comprises several scenarios with real-world traffic data and different features. To this end, the setup relies on the extensive collection of traffic flow readings made public by the Madrid City Council in its Open Data portal (<https://datos.madrid.es/>), which releases a 5-year historic record of 15-minute traffic flow observations collected in more than 3,800 locations over this city. With these data, together with meteorological and



Fig. 4. Distribution of the selected traffic loops located in the city center of Madrid (Spain), whose collected data have been used for the experiments reported in this work.

calendar information, different datasets \mathcal{X} have been composed. Hence, an scenario is defined based on the datasets in use, the considered forecasting horizon (i.e., the gap between the query time and the time at which the target variable sought occurs), and the forecasting model used to model the relationship between its input variables and the traffic measure to be predicted. By examining the relative behavior of the confidence metrics in each one of such scenarios, we can 1) gauge how each dimension of the scenario affects uncertainty (in response to RQ2), and 2) compare uncertainty estimation methods in terms of the metrics in different contexts (providing informed insights for RQ1).

A. Datasets

The data used for defining the scenarios introduced above are based on real urban traffic data collected in 10 different locations of the city of Madrid (Spain). The selection of inductive loops has been made according to the variability of the surrounding urban topology of every sensor, in order to analyze the potential relationship between the traffic profiles in those locations and the measured uncertainty of a modeling pipeline aimed to forecasts those traffic measures. The spatial distribution of the selected sensors, depicted in Figure 4, shows that some loops are placed in locations with heavy traffic conditions (i.e., references¹ 3697, 3910 and 5761), while others are placed in small roads of residential areas or main arterial roads traversing commercial districts.

Data collected for the whole year 2019 were built into datasets to train and test different forecasting models. Such datasets include information of up to 5 steps (i.e., 1 hour and 15 minutes) of past traffic measurements before the traffic point to be predicted (t_{-4} , t_{-3} , t_{-2} , t_{-1} , and t_0). This

¹These numerical references correspond to the labels assigned to the loops in the Open Data portal from which data were retrieved.

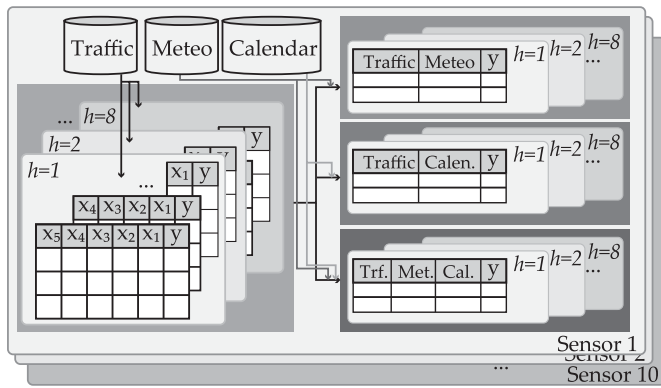


Fig. 5. Datasets created for the experiments.

window of observations will be configured as the input to the models in 2 forms $\omega = \{1, 5\}$: either all past measurements ($\omega = 5$) or just one ($\omega = 1$). For each case, different forecasting horizons $h = \{1, 2, 4, 8\}$ are considered, producing predictions in instants t_{+1} , t_{+2} , t_{+4} and t_{+8} , and providing traffic forecasts of up to two hours in the future. Datasets may include (or not) meteorological information (temperature, cloud cover, humidity and precipitation intensity) and calendar information regarding local and national holidays, academic and scholar calendar and day of the week. The availability of this data at the input to the models is denoted by two binary variables $m = \{0, 1\}$ (meteorology) and $c = \{0, 1\}$ (calendar). With these different types of input variables, datasets with all possible combinations are created following the scheme in Figure 5, leading to a collection of 320 datasets resulting from 10 loops \times 2 window lengths $|\omega| \times |m|$ (meteorological information available or not) $\times |c|$ (calendar information available or not) \times 4 forecasting horizon values. Models will be tested and their uncertainty measured with different sets of these combinations, in order to extract insights and answer each of the research questions stated in the introduction.

Additionally, a standardization procedure was applied to all data to guarantee the correct operation of some of the models, particularly those that are sensitive to differences in the statistical support of their predictors. Lastly, datasets were split into train and test partitions, stratifying them across all months of the year available for training. For each month, 3 weeks were considered as train data, whereas 1 week was left for test data. This permits to include all kinds of traffic profiles (which are highly variable throughout the year) in the training dataset, and also allows for a proper train-calibration-test split for the Conformal Prediction approach: the calibration set will consist of the 3rd week of each month, thus providing data from all along the year to calibrate the models. Techniques that do not include a calibration step do not use this part of the dataset as training data, so all the models receive the same training information.

B. Methodological Approach

Departing from the collection of 320 datasets defined in the previous section, a processing pipeline is established. The steps are schematically depicted in Figure 6.

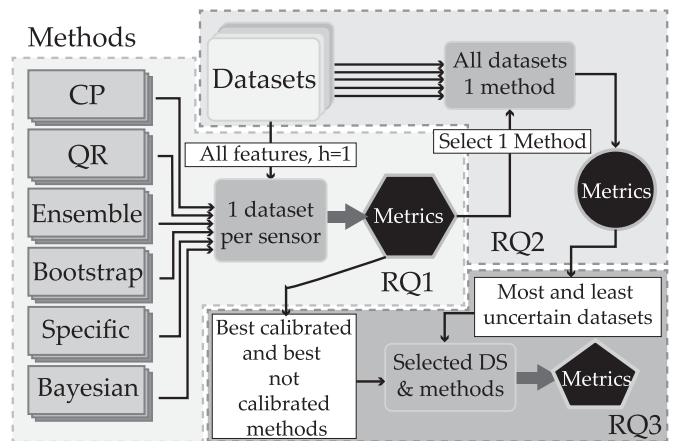


Fig. 6. Schematic diagram of the processing pipeline applied to the traffic forecasting datasets under analysis.

To answer RQ1, several predictive methods that cover all the uncertainty estimation approaches described in Section II-C are considered. Prediction tests are conducted with each model and the dataset combination comprising all features ($m = c = 1$) and the forecasting horizon set to $h = 1$. The uncertainty of each (model, dataset) combination is then estimated with the available techniques. On the other hand, the predictive performance of the model has been measured by means of the coefficient of determination R^2 . This score is bounded between 0 and 1, and as opposed to other metrics such as RMSE, allows for a meaningful comparison between datasets that may have different numerical ranges. A relevant remark related to this first comparison among techniques is that not all techniques are suitable for all algorithms. For instance, measuring uncertainty through techniques closely linked to specific elements of deep neural networks (e.g., Monte Carlo dropout) are not available for tree-based shallow learning algorithms. Thus, the comparison of techniques will be influenced by the capacities and limitations of the underlying predictive models, as well as by the *epistemic uncertainty* they introduce. Table I shows the specific models that have been considered for each uncertainty estimation technique in our experimental benchmark. We note that the time granularity of the datasets in use (15-minute slot) and the traffic variable under target (traffic flow in vehicles per unit of time) do not suggest that spatial relationships between loops can be exploited towards a better forecasting performance. This statement does not hold in scenarios considering other traffic variables collected at higher rates in simpler road networks, such as speed in highways. Therefore, graph-based models (e.g. GCN) are not considered in our experiments.

In addition to the predictive performance of the considered forecasting models (which may differ among them), the comparison benchmark also accounts with the confidence metrics described in Section II-D. By simultaneously reporting on both confidence and forecasting error, we will be able to compare them in model-agnostic dimensions, connecting clearly with the responses sought to answer RQ1. When it comes to RQ2, a single uncertainty estimation method and predictive model will be selected, based on the results

TABLE I
PREDICTIVE MODELS USED IN THE CASE STUDY AND THE UNCERTAINTY TECHNIQUES APPLICABLE TO EACH OF THEM

Model	Conformal Prediction	Ensemble	Quantile Regression	Bootstrapping	Bayesian	Specific	Details
RFR	CP-RFR	E-RFR	×	B-RFR	×	×	100 estimators
ETR	CP-ETR	E-ETR	×	B-ETR	×	×	100 estimators
GBR	CP-GBR	×	Q-GBR	B-GBR	×	×	loss: quantile, 100 estimators
ABR	CP-ABR	E-ABR	×	B-ABR	×	×	100 estimators, learning rate: 0.1
MLP	CP-MLP	×	×	B-MLP	BNN	×	Dense(50), 20 epochs, Adam
DL	×	E-DL	QR-DL	B-DL	×	EvDL, MCD, HR, MIS	LSTM(50, ReLu), Dense(20, ReLu), 20 epochs, Adam

RFR: Random Forest Regressor; ETR: Extra Trees Regressor; GBR: Gradient Boosting Regressor; ABR: Ada Boost Regressor; MLP: Multilayer Perceptron; DL: Deep Learning; QR-DL: Deep Learning with Quantile Loss; BNN: Bayesian Neural Network; B-*: Bootstrapping; HR: Heteroscedastic Regression; EvDL: Evidential Deep Learning; MCD: MonteCarlo Dropout, MIS: Mean Interval Score Regression




Fig. 7. Average flow in each location (red line). The shaded area represents \pm one standard deviation.

obtained from the prior experimentation done in regards to RQ1. The rest of dataset combinations (different horizons and different sets of features) will be used to obtain forecasts, estimate uncertainty and assess their impact on the confidence levels. Lastly, tests with calibration-based and non-calibrated techniques will be conducted using different combinations of datasets, so that RQ3 can be replied in an empirically informed fashion.

IV. RESULTS AND DISCUSSION

We now discuss on the experimental results of the setup described in Figure 6. In doing so, we organize our examination of such results in terms of the research questions posed in the introduction. Models and uncertainty estimation techniques have been tested with different sensors placed in around the city in order to have different traffic profiles and also assess the impact of the traffic behavior on the modeling tools. Thus, traffic profiles of the 10 traffic sensors are shown on Figure 7. Each plot nested in this figure represents the average traffic (computed over 96 daily traffic traces) of every sensor in the

i^{th} 15-minute interval of a day, where the standard deviation is also provided as a shaded area to illustrate the intra-day variability of the traffic profiles at each sensor’s location.

As revealed in this figure, sensors such as 3642, 6132 or 3500 have very low traffic profiles, with less than 300 vehicles/hour on average for the busiest hour. This implies more real traffic data close to 0 vehicles/hour and larger relative errors in the prediction. Large deviation areas like the one in sensor 4458 anticipate larger confidence intervals, while loops like 4192 with narrower deviations and wider dynamic range are expected to provide better performing and less uncertain models. As a sample of how intervals are produced, a single day of data collected by sensor 4458 is shown on Figure 8, in which the real data and the confidence intervals produced by the combination of conformal prediction (CP) and a Random Forest regressor (RFR) are presented. The interval is highly informative, as its width unveils when the most uncertain predictions are produced, and how uncertain they may be. A prediction at noon with a 90% confidence will be quite uncertain.

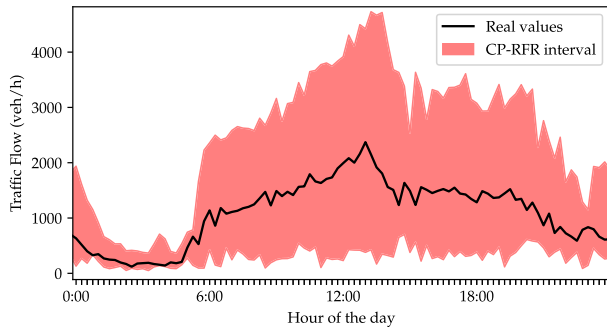


Fig. 8. One day of traffic data collected in loop 4458. Real values are shown on top of the confidence intervals estimated by conformal prediction for a random forest forecasting model.

Once the inherent variability of traffic data has been shown, we proceed by discussing on the quantitative results produced for addressing every research question:

A. RQ1: How Do Different Uncertainty Quantification Techniques Perform When Applied to a Particular Traffic Forecasting Model?

The first research question regards the performance of the uncertainty estimation techniques considered in our benchmark. As indicated in Table I, not all of such techniques operate in the same fashion, reason for which the metrics proposed in Section II-D are used as a standardized framework that allows comparing them under unified criteria.

We start our discussion around RQ1 by inspecting Figure 9, which shows the mean interval length, or the average width of the interval (MIL) obtained for the datasets comprising meteorological and calendar features ($m = c = 1$), using each predictive model and uncertainty estimation approach. On top of each bar plot, the performance score R^2 is shown with a red line. As expected, performance is generally lower for those sensors with lower traffic profiles, particularly for 3500, which scores the worst with R^2 values not surpassing beyond 0.7. The lower performance of ADABOOST based methods (namely, CP-ABR, E-ABR and B-ABR) is specially visible here, but can be also noticed in the rest of sensors. This predictive method suffers the burden of removing data or removing estimators for estimating uncertainty, and in general is not a suitable option if confidence is to be measured. On the other hand, deep learning approaches (including heteroscedastic regression) present generally wider intervals than conformal prediction approaches except in the case of MIS regression, which stands closer to QR-DL. A noteworthy increase is perceived when using evidential formulations of deep learning models. Indeed, the confidence intervals provided by EvDL are significantly wider than those elicited by other uncertainty estimation techniques, except for sensor 5761. This can be explained by the renowned sensitivity of evidential formulations with respect to its regularization coefficient λ [88]. This value was left equal to a constant value (1.0) for all experiments so as to expose this known issue. Oppositely, inspecting the behavior of bootstrapping methods reveals a similar performance to other model-agnostic techniques in terms of predictive capacity and MIL, yet requiring

a significantly higher computational effort to estimate its prediction intervals (one model must be learned for each of the 100 bootstraps configured for the experiment).

On the other hand, conformal prediction (CP) based methods perform very similarly to each other across different traffic loops in terms of the MIL metric. This holds in all cases except for CP-ETR in two of the sensors (10124 and 3500). The removal of certain nodes that does not happen in the very similar Random Forest model produces extreme peaks in the estimation of the interval boundaries when real observations are very close or equal to 0. This explains these high values that are produced after averaging these peaks with the rest of the points of the boundary, which are essentially almost the same as for Random Forest. Conformal prediction, quantile regression, bootstrapping or ensemble methods produce intervals of very similar widths for most of the cases, while Bayesian methods present an irregular behavior in terms of interval width and a subpar predictive performance reflected in generally lower R^2 values.

Measuring the average interval width through the lens of the MIL metric shows one face of the uncertainty associated to a traffic forecast: the narrower the interval, the more trustworthy the prediction can be thought to be. However, too narrow intervals can leave too many real samples outside them. This reason motivates the adoption of ICP as the second confidence metric for the discussion of this first set of results.

Figure 10 depicts the distributions of the interval coverage ICP of each method over the different sensors. The red line represents the statistical significance α established for the confidence levels. In principle, the coverage of all methods should be close to this line, disregarding the input dataset, for the sake of a reliable behavior. It is noticeable that CP-based methods achieve a stable and close-to- α distribution of the ICP metric: this implies that the interval width values discussed before steadily cover 90% of the real traffic samples, rendering them useful to know the model uncertainty. Quantile approaches have a larger dispersion, specially for deep learning, which can also be perceived for MIS regression, closer to quantile approaches than to techniques like dropout or evidential learning. This means that for intervals with similar width, the amount of real samples that are left out the interval is more dependent on the input data. Results for obtaining uncertainty estimates from the ensemble version of AdaBoost are clearly the worst, with confidence intervals similar in size to others, but covering much less real data (they are wider in some areas but much narrower in others, providing similar average widths, but worse coverage). This is linked to the way in which AdaBoost builds each estimator using information from the previous ones: since it is a sequential process in which every estimator specializes in difficult-to-forecast examples, weak learners composing the boosting ensemble are not by themselves reliable forecasting models [108]. Lastly, we again observe that deep learning based approaches (also HR and BNN) produce larger confidence intervals. The coverage analysis of these wider intervals reveals that, naturally, they cover more real samples. For evidential deep learning, very large intervals cover up to the entirety of the real data. This could appear as convenient, as the intervals cover more




Fig. 9. Mean Interval Length (MIL) obtained by the uncertainty estimation methods over different forecasting models learned from the data collected by every loop. The forecasting horizon is $h = 1$. Predictive performance is presented in red (R^2 score).




Fig. 10. Interval Coverage Percentage (ICP) obtained by the uncertainty estimation methods over different forecasting models learned from the data collected by every loop. Boxplots represent the distribution of ICP values over the 10 loops with a forecasting horizon $h = 1$. Red line represents the established confidence level $\alpha = 0.9$.

real samples, but it is far from informative and actionable, as the intervals do not comply with the established significance level: the desired output is the narrower interval that covers the established amount of samples. This can be appreciated in Figure 11, where intervals produced for two test days in loop 4458 for BNN, EvDL and CP-RFR are compared to each other. Predictive performances appear to be very similar: forecasts of BNN, in purple, show a higher deviation with




Fig. 11. Two days of test traffic data collected in loop 4458. Real values (ground truth) compared to forecasts given by three differently performing methods, together with their estimated confidence intervals for $\alpha = 0.9$.

respect to their ground truth at some points, yielding a slightly lower R^2 score. The intervals of EvDL appear to be similar for a large part of the day, but they are overestimated in the night periods, leading to less useful confidence information for actionability purposes (e.g. adaptive street lighting in the concerned part of the road network). In the case of BNN, its confidence interval seems to be more similar to that produced by CP-RFR in this particular scenario, but it tends to be wider on the upper boundary, leaving space for more samples, and thereby increasing the value of the ICP metric above its target confidence level α .

On the other side, approaches featuring larger interval widths also present a higher dispersion in their measured ICP values, revealing their higher dependency on the input data. In general, CP-based approaches show the most stable

TABLE II
CONFIDENCE METRICS OF CP-RFR OVER DIFFERENT DATASET AND TRAFFIC SENSORS

Sensor	Metric	All Features ($m = c = 1$)				Traffic + Calendar ($m = 0, c = 1$)				Traffic + Metro ($m = 1, c = 0$)				Only Traffic ($m = c = 0$)			
		$h=1$	$h=2$	$h=4$	$h=8$	$h=1$	$h=2$	$h=4$	$h=8$	$h=1$	$h=2$	$h=4$	$h=8$	$h=1$	$h=2$	$h=4$	$h=8$
4458	R^2	0.94	0.92	0.86	0.75	0.94	0.92	0.86	0.75	0.94	0.91	0.85	0.71	0.94	0.91	0.84	0.69
	ICP	0.89	0.89	0.88	0.88	0.89	0.89	0.89	0.89	0.89	0.89	0.89	0.89	0.89	0.89	0.89	0.89
	MIL	527.76	640.15	850.82	1220.51	531.66	646.57	865.51	1201.72	540.41	676.08	921.57	1328.06	557.76	683.74	925.55	1292.41
6980	R^2	0.84	0.81	0.75	0.61	0.84	0.81	0.74	0.61	0.84	0.8	0.72	0.56	0.83	0.79	0.7	0.53
	ICP	0.91	0.9	0.91	0.9	0.9	0.9	0.91	0.91	0.91	0.9	0.91	0.91	0.9	0.9	0.91	0.91
	MIL	216.62	231.81	275.49	344.23	218.71	230.96	279.92	352.82	218.93	241.38	289.64	378.44	223.22	240.9	293.27	380.8
10124	R^2	0.79	0.77	0.7	0.57	0.78	0.76	0.69	0.56	0.78	0.76	0.66	0.49	0.77	0.73	0.64	0.45
	ICP	0.9	0.91	0.9	0.89	0.9	0.9	0.9	0.9	0.9	0.91	0.91	0.9	0.9	0.9	0.9	0.89
	MIL	137.79	146.69	174.66	210.66	136.62	148.33	174.03	214.51	138.46	154.19	189.33	232.42	143.1	157.88	188.92	232.06
6132	R^2	0.78	0.76	0.69	0.52	0.78	0.75	0.68	0.5	0.77	0.74	0.65	0.45	0.76	0.73	0.63	0.42
	ICP	0.91	0.9	0.9	0.9	0.91	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.91	0.91	0.9	0.9
	MIL	93.94	97.12	106.7	133.95	95.64	99.38	107.25	133.89	95.65	99.52	114.34	140.92	98.57	103.42	115.93	147.53
3642	R^2	0.87	0.84	0.75	0.63	0.87	0.83	0.74	0.63	0.86	0.82	0.71	0.56	0.86	0.81	0.68	0.53
	ICP	0.89	0.9	0.89	0.9	0.9	0.9	0.89	0.9	0.9	0.9	0.89	0.91	0.9	0.9	0.88	0.89
	MIL	107.04	121.9	140.59	187.3	109.15	124.02	142.84	185.65	111.11	126.93	152.12	207.48	112.9	130.03	155.45	196.81
4192	R^2	0.95	0.93	0.89	0.8	0.95	0.93	0.89	0.81	0.95	0.93	0.87	0.75	0.94	0.92	0.87	0.75
	ICP	0.91	0.91	0.91	0.91	0.91	0.91	0.91	0.91	0.91	0.91	0.91	0.91	0.92	0.91	0.91	0.91
	MIL	466.4	541.57	692.87	983.08	475.68	544.31	700.7	967.34	484.22	565.85	741.42	1055.89	502.96	574.93	757.04	1084.64
3697	R^2	0.95	0.94	0.9	0.81	0.95	0.94	0.9	0.8	0.95	0.93	0.89	0.79	0.95	0.93	0.89	0.78
	ICP	0.91	0.9	0.9	0.91	0.9	0.9	0.9	0.9	0.91	0.9	0.91	0.91	0.9	0.9	0.9	0.92
	MIL	473.44	544.55	670.17	959.55	459.33	547.35	677.76	959.58	489.7	565.44	731.49	1029.71	485.88	567.11	728.05	1055.73
3910	R^2	0.95	0.93	0.88	0.78	0.95	0.93	0.88	0.79	0.95	0.92	0.87	0.73	0.94	0.92	0.86	0.72
	ICP	0.91	0.9	0.91	0.9	0.91	0.9	0.9	0.9	0.91	0.9	0.9	0.9	0.9	0.9	0.9	0.91
	MIL	283.58	331.51	428.51	626.99	285.29	325.0	419.52	595.69	293.95	347.2	453.5	660.41	296.98	348.3	454.07	657.75
3500	R^2	0.64	0.62	0.56	0.45	0.63	0.61	0.55	0.47	0.63	0.6	0.53	0.41	0.62	0.59	0.52	0.41
	ICP	0.89	0.9	0.89	0.9	0.9	0.9	0.89	0.9	0.9	0.9	0.89	0.9	0.9	0.89	0.9	0.9
	MIL	70.49	74.86	79.94	94.59	72.6	75.89	79.99	93.54	73.59	76.45	82.58	100.66	73.96	76.03	86.1	99.97
5761	R^2	0.94	0.92	0.88	0.78	0.94	0.92	0.87	0.77	0.94	0.91	0.86	0.74	0.93	0.91	0.86	0.73
	ICP	0.91	0.91	0.9	0.9	0.91	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.91
	MIL	725.49	843.83	1032.46	1369.65	722.01	835.	1043.94	1417.15	727.19	849.99	1063.37	1441.41	753.34	854.84	1098.85	1474.63

behavior independently of data and underlying model, due to the calibration process whose importance will be analyzed in subsequent sections.

B. RQ2: Which Changes in the Available Data Affect Uncertainty? How Can the Confidence of a Forecasting Model Be Improved by Changing/Augmenting Data at Its Input?

After exploring a general performance view of all proposed methodological approaches in all available locations, the best performing model and uncertainty estimation technique is chosen in order to delve into the way in which the configuration of the input dataset \mathcal{X} affects the uncertainty of forecasts issued by such models. We select the CP-RFR combination (namely, Conformal Prediction with Random Forest Regressor), so that both ICP and MIL metrics are computed over the different dataset configurations resulting from the reduction of the number of input features and the consideration of different prediction horizons h . Before proceeding with the analysis of the obtained results, we emphasize that the aim of this research question is not to examine whether uncertainty estimation can be an effective criterion for feature selection. Instead, RQ2 aims to expose that increments in performance due to the inclusion of relevant features (which might contribute to the model's performance, and could hence be identified by any given feature selection method) and/or the removal of samples can be counterproductive when it comes to the uncertainty at the output of the model.

Table II summarizes the metrics obtained in all these experiments. Several clear patterns arise in the light of these

results. To begin with, an expected degradation of predictive performance occurs for all situations when the predictive horizon h is increased. However, removing features from the original complete dataset has, in most cases, a slim impact in what regards to the prediction error: forecasting with all features renders very similar predictive scores that considering only traffic and calendar (i.e., disregarding meteorological features). There are cases in which removing the meteorological features produces better results (sensors 3500 and 3910) for $h = 4$, from where one may conclude that these features do not contribute significantly to the model's performance in this circumstances. Analogously, when traffic and meteorological datasets are fed, the differences to the case when only traffic data is used for forecasting are negligible in most cases (with exceptions like the higher predictive horizons for sensor 10124). This minor contribution of weather features to the model's performance was identified in [109], as a consequence of an highly stable weather in this city, and the inherent relation of certain weather conditions with calendar features. This conclusion should motivate the community to go beyond predictive performance and also examine whether additional features contribute anyhow to the reduction of the uncertainty associated to the forecasts.

Secondly, when the focus is set on the ICP metric, its steady behaviour throughout the whole set of scenarios is remarkable albeit expected, as the calibration stage of Conformal Prediction is precisely aimed at achieving this. The next section will revolve around this in connection to RQ3, analyzing the consequences of obtaining confidence intervals without any calibration. Furthermore, considering that the values of ICP guarantee that the demanded 90% of true samples are

TABLE III
PERFORMANCE AND INTERVAL WIDTH WITH COMPLETE AND ABLATED DATASETS

Metric	Dataset	4458	6980	10124	6132	3642	4192	3697	3910	3500	5761
R^2	Complete	0.94	0.83	0.77	0.76	0.86	0.94	0.95	0.94	0.62	0.93
	25% Ablated	0.94	0.83	0.77	0.77	0.85	0.94	0.95	0.94	0.63	0.93
	50% Ablated	0.93	0.85	0.77	0.74	0.84	0.93	0.94	0.92	0.64	0.93
	75% Ablated	0.91	0.84	0.77	0.76	0.84	0.93	0.94	0.93	0.64	0.93
MIL	Complete	525.14	216.56	137.6	94.11	105.98	463.15	484.05	280.28	72.52	730.58
	25% Ablated	535.12	216.66	138.9	98.64	106.78	433.78	489.04	281.26	67.39	751.18
	50% Ablated	570.21	218.2	142.62	99.07	109.58	450.96	501.2	296.05	75.96	763.94
	75% Ablated	584.73	223.79	147.09	100.68	113.18	472.89	509.39	305.67	80.35	769.93

inside the intervals provided in all cases, the width of such intervals can be compared fairly to each other. In general, trends noted in the MIL metric are similar to those found in performance measurements: the more difficult to predict, the wider the interval. However, this measurement provides another perspective of the contribution of different variables in each case. For instance, increasing the prediction horizon from $h = 1$ to $h = 8$ results in a R^2 degradation of around 0.2 for loops 4458 and 3500. For this similar performance degradation, the confidence interval is 2.27 times wider in $h = 8$ than in $h = 1$ for loop 4458, while it is only 1.3 times wider in the other sensor. Thus, performance and confidence intervals do not degrade equally, because they are not co-linear. This means that performance itself is not enough to estimate uncertainty. On the other hand, it is possible to observe some situations in which removing variables favors the confidence; in some of them like in sensor 6980 with forecasting horizon $h = 2$ the change is small, but the model is more confident without meteorological features: Traffic and calendar presents a narrower interval than all features, and only traffic is also narrower than traffic and meteorological features. This seems an interesting way of analyzing whether a set of features do contribute to the confidence of the model. This happens for loop 4458, where weather features have a positive impact in both cases. Moreover, there are cases in which removing a variable results in narrower intervals, such as loop 3697 with meteorological features and only traffic features. This improvement in uncertainty could be due to statistical variance of these particular data (MIL is computed by averaging over thousands of samples), but in any case it reinforces the intuition that these variables have a meager contribution to performance.

Lastly, we conduct a data ablation study to grasp a better understanding of the impact of training data on the levels of uncertainty estimated at the output of traffic forecasting models. By systematically removing data points from the training datasets, we can analyze how these ablations affect both the performance and the measured uncertainty of the models. Experiments designed for this purpose consist of 1) removing a percentage of the training and calibration data, and 2) comparing the uncertainty metrics to those elicited by models learned from the complete data. In doing so, the rest of the considered study variables (types of features, horizon, modelling approach) are fixed, whereas the ablation of data is performed over the traffic datasets with prediction horizon $h = 1$. We follow the scheme described in Section III-A, i.e., training and calibration data are spread throughout the whole

year, with portions of data reserved for test to guarantee that all seasonal variations are captured by the models. Thus, every 4 weeks there is a block of 2 weeks of training data, 1 week of calibration data and 1 week of test data. For this experiment, the first 3 weeks of a subset of these blocks are removed. Then, models are trained without these data and tested again. The experiment has been repeated with 3 different levels of ablation, removing 25%, a 50% and a 75% of the training and calibration data.

Results of this set of experiments are summarized in Table III. These results reveal that for all loops, the predictive performance given by the R^2 metric reported therein remains almost unaltered when removing training data. This negligible impact of reducing the training data in the predictive capabilities of traffic forecasting models was expected as per the conclusions of the study in [110], which showed that the models can be trained with very reduced datasets, achieving similar levels of performance to those with larger training datasets, due to the stability in traffic data patterns for those particular locations. However, when it comes to uncertainty, it can be observed that interval sizes are increased for most sensors, slightly when the percentage of removed data is small, but considerably as this percentage is increased. Considering this last group, the increase of interval sizes ranges between 5% to 11%, which can be attributed to the lack of some particular kinds of days in the training data. For instance, the whole data from the month of August are removed from training and calibration, but there is a test week with August days, whose traffic usually features a very different statistical behavior than the rest of the months. While forecasting errors are smoothed due to the strong auto-correlation existing between the target variable and the features (behavior changes in test week will also be reflected in the features), the confidence metrics reflect a general increase due to the absence of these data in the calibration tests. It is interesting to observe that the sensors with least relative change (namely 4192, 3697 and 5791) are those with narrower variances (Figure 7).

When confidence estimation techniques are properly calibrated, the size of the confidence interval can be a valuable indicator of the way in which features contribute to the confidence levels of the output. Thus, besides obtaining an actionable piece of information with the confidence intervals that guarantee the inclusion of a certain amount of true samples, the analysis of these intervals allows for further insights than can be helpful when designing the dataset \mathcal{X} for learning the forecasting model. In the particular datasets considered in this study, it is apparent that meteorological

TABLE IV
CONFIDENCE METRICS OF CALIBRATED AND NON-CALIBRATED APPROACHES OVER SCENARIOS WITH ONLY TRAFFIC DATA ($m = c = 0$)

Sensor	Metric	CP-RFR				E-RFR				HR				EvDL			
		$h=1$	$h=2$	$h=4$	$h=8$	$h=1$	$h=2$	$h=4$	$h=8$	$h=1$	$h=2$	$h=4$	$h=8$	$h=1$	$h=2$	$h=4$	$h=8$
4458	R^2	0.94	0.91	0.84	0.69	0.94	0.91	0.84	0.69	0.94	0.91	0.83	0.64	0.91	0.87	0.77	0.47
	ICP	0.89	0.89	0.89	0.89	0.84	0.85	0.85	0.83	0.98	0.97	0.98	0.99	0.98	0.93	0.74	0.23
	MIL	557.76	683.74	925.55	1292.41	450.69	545.88	735.34	1015.81	709.93	874.9	1239.66	2030.38	1462.95	940.43	573.01	180.88
6980	R^2	0.83	0.79	0.7	0.53	0.83	0.79	0.7	0.53	0.83	0.79	0.7	0.53	0.79	0.72	0.56	0.31
	ICP	0.9	0.9	0.91	0.91	0.86	0.87	0.87	0.87	0.99	0.98	0.98	0.99	1.0	1.0	1.0	0.99
	MIL	223.22	240.9	293.27	380.8	179.96	202.18	243.01	301.56	292.54	332.53	403.14	520.31	1564.44	1272.89	948.28	915.4
10124	R^2	0.77	0.73	0.64	0.45	0.77	0.73	0.64	0.46	0.78	0.73	0.62	0.45	0.74	0.68	0.53	0.27
	ICP	0.9	0.9	0.9	0.89	0.86	0.86	0.86	0.86	0.96	0.96	0.96	0.97	1.0	0.99	0.96	0.81
	MIL	143.1	157.88	188.92	232.06	115.38	126.85	146.69	176.21	155.12	176.96	218.54	300.66	488.73	378.8	312.39	165.98
6132	R^2	0.76	0.73	0.63	0.42	0.76	0.73	0.63	0.42	0.77	0.74	0.64	0.45	0.71	0.67	0.49	0.33
	ICP	0.91	0.91	0.9	0.9	0.86	0.86	0.87	0.87	0.97	0.97	0.97	0.97	0.99	0.98	0.89	0.49
	MIL	98.57	103.42	115.93	147.53	76.42	80.95	92.79	113.56	103.83	112.81	134.65	183.4	164.28	155.25	110.06	50.95
3642	R^2	0.86	0.81	0.68	0.53	0.85	0.81	0.68	0.53	0.84	0.78	0.66	0.5	0.83	0.75	0.59	0.32
	ICP	0.9	0.9	0.88	0.89	0.86	0.86	0.86	0.85	0.97	0.97	0.98	0.98	0.98	0.97	0.82	0.55
	MIL	112.9	130.03	155.45	196.81	90.19	103.11	127.87	157.05	134.26	162.31	216.65	291.06	218.23	208.17	120.27	76.29
4192	R^2	0.94	0.92	0.87	0.75	0.94	0.92	0.87	0.75	0.94	0.91	0.85	0.71	0.92	0.89	0.78	0.61
	ICP	0.92	0.91	0.91	0.91	0.85	0.85	0.85	0.85	0.96	0.97	0.97	0.98	1.0	0.99	0.58	0.35
	MIL	502.96	574.93	757.04	1084.64	375.39	441.29	569.86	775.43	585.99	727.23	995.07	1460.43	2831.65	772.97	357.99	298.92
3697	R^2	0.95	0.93	0.89	0.78	0.95	0.93	0.89	0.78	0.95	0.93	0.87	0.75	0.93	0.89	0.77	0.59
	ICP	0.9	0.9	0.9	0.92	0.85	0.85	0.85	0.85	0.98	0.98	0.98	0.97	0.95	0.89	0.4	0.2
	MIL	485.88	567.11	728.05	1055.73	380.29	447.0	569.21	789.45	646.82	765.92	992.71	1391.19	720.45	686.63	228.62	142.76
3910	R^2	0.94	0.92	0.86	0.72	0.94	0.92	0.86	0.72	0.94	0.91	0.84	0.67	0.91	0.86	0.63	0.38
	ICP	0.9	0.9	0.9	0.91	0.86	0.85	0.86	0.86	0.97	0.97	0.97	0.97	0.99	0.97	0.53	0.67
	MIL	296.98	348.3	454.07	657.75	232.68	276.74	362.88	505.82	345.72	424.19	716.0	961.62	852.71	726.62	192.17	368.8
3500	R^2	0.62	0.59	0.52	0.41	0.62	0.59	0.52	0.41	0.65	0.62	0.56	0.44	0.61	0.6	0.52	0.38
	ICP	0.9	0.89	0.9	0.9	0.86	0.86	0.86	0.86	0.97	0.97	0.97	0.97	1.0	1.0	1.0	1.0
	MIL	73.96	76.03	86.1	99.97	60.47	64.93	72.12	81.35	96.12	103.35	115.55	135.03	880.21	927.87	769.13	530.5
5761	R^2	0.93	0.91	0.86	0.73	0.93	0.91	0.86	0.73	0.93	0.91	0.85	0.7	0.91	0.89	0.76	0.61
	ICP	0.9	0.9	0.9	0.91	0.85	0.85	0.85	0.85	0.97	0.97	0.97	0.98	0.55	0.71	0.18	0.41
	MIL	753.34	854.84	1098.85	1474.63	581.63	675.8	855.98	1151.84	895.51	1044.98	1434.81	2144.94	357.83	551.66	132.8	480.21

features contribute slightly in terms of predictive performance, but having the intervals not only confirms this point, but it even stresses it out for some sensors. For example, loop 5761 has equal results in terms of performance for all features and for traffic and calendar datasets, but the latter presents smaller confidence intervals. This is an interesting outcome, as it shows that in the case of slight variance in the datasets, like the case of stable meteorological data, models can be unable to represent small variations, but estimating uncertainty can reveal these nuances.

C. RQ3: What Impact Does the Calibration Process of Some of the Uncertainty Estimation Techniques Have on Their Outcome? Why Is It Relevant for Traffic Forecasting?

The last research question relates to the calibration of models and its impact on the quality of the estimated uncertainty. Certainly, the concept of *calibration* is not new [111], but is relatively overseen in a broad part of the literature focused on uncertainty estimation. The calibration process aims at guaranteeing the properties of the output interval. In the results shown in previous section it is possible to observe that regardless the dataset attributes, if the confidence estimation technique is calibrated, the output interval always covers a percentage of true samples very close to the significance level established by design. This means that the width of the intervals can be trusted (an important feature for the trustworthiness of the uncertainty estimation and the actionability of forecasts) and always have the same meaning among the considered combinations of models and estimation techniques.

In order to assess the relevance of this calibration process, we note that it is an essential part of Conformal Prediction.

Therefore, we perform a comparison of the aforementioned CP-RFR to a non-calibrated uncertainty estimation counterpart (E-RFR), as well as to other methods characterized by an unstable ICP behavior in Figure 10, namely, HR and EvDL. The outcomes of this comparison are presented in Table IV, considering the least favorable datasets (only traffic) as the effects of the lack of calibration on the estimated uncertainty become more noticeable with less variability in the input. As the calibration process consists precisely of statistically characterizing the behavior of such a variability, in low-variability scenarios out of distribution data instances may have more weight in the definition of the intervals. Uncertainty estimation techniques that do not consider this can end up providing 1) wider intervals to cover more real samples than the specified percentage (due to the fact that they consider for the intervals points that are outside the significance boundaries in the training data distribution); or 2) narrower intervals that do not meet the statistical significance specifications.

The first 4 columns of Table IV correspond to the calibrated method: in them, it is possible to observe the persistence of the ICP value for all loops and all values of the forecasting horizon h . When considering the same forecasting model (RFR) but using the ensemble technique to estimate the intervals, very similar R^2 scores are obtained (as expected, as they essentially resort to the same algorithm for producing the forecasts). However, in this case the whole set of results appears to be under-calibrated, as the coverage of the estimated intervals is lower than the expected one (85%), giving place to narrower intervals. Should these intervals and performance levels be considered to compare them to those elicited by CP-RFR, the comparison would declare E-RFR as the best forecasting method. Nonetheless, it is leaving 5% of the




Fig. 12. Calibration curves (colored in blue) corresponding to each case in Table V. The dashed orange line corresponds to the ideal case where the confidence intervals estimated for the model are perfectly calibrated. The area between the calibration curve and the ideal case is given in each figure.

real samples outside the coverage of its confidence intervals, thereby penalizing the trustworthiness of the estimated uncertainty. HR seems to work the other way around, by increasing the size of the interval up to covering almost all real cases (ICP close to 99%), and obtaining similar regression results. Lastly, EvDL was shown in the previous section that it failed to provide reliable confidence levels with the complete dataset and horizon $h = 1$: the reported interval width tended to be wider and ICP values were close to 1 in most cases. When only-traffic data were provided, this method performed erratically. Apparently, wider intervals and ICP values close to 1 are consistent with $h = 1$, but while for some sensors this behavior is sustained, others present a notable decay in predictive performance (see the cases of sensors 4458, 6980, 3910, in comparison to other methods), and also in the size of the interval, reducing the coverage to very low levels (even to only 20% of samples in loop 3697). This demeanor is sharper for those sensors with larger dynamic ranges, suggesting the susceptibility of the method to deal with data in the tails of the distribution unless 1) its regularization parameter λ is properly tuned; or 2) by performing a posterior calibration of the model's output.

Additional tests were conducted considering CP-RFR and levels of confidence equal to $\alpha = 0.85$ and $\alpha = 0.99$ in order to assess if such methods are just providing the intervals that are proper for other levels of confidence. Table V reports, for one of the sensors, a trend that prevails over the totality of loops analyzed in these experiments: in general, intervals obtained for $\alpha = 0.9$ that cover 85% of the real traffic samples are narrower than confidence intervals provided by CP-RFR for $\alpha = 0.85$. The same behavior occurs when $\alpha = 0.99$. One may arrive at the conclusion that intervals computed by these methods for $\alpha = 0.9$ are better to estimate intervals with support equal to 0.85 and 0.99, as they provide narrower results; however, and unlike with CP, these results lack any statistical guarantees.

To further argue on the need for calibrated traffic forecasts, we inspect the relationship between different confidence levels α and the observed traffic values for the forecasting horizon values and uncertainty estimation techniques reported

TABLE V
CONFIDENCE METRICS FOR CP-RFR WITH DIFFERENT CONFIDENCE LEVELS α COMPARED TO OTHER METHODS WITH A FIXED CONFIDENCE LEVEL

Sensor 5761		Metric			
		$h=1$	$h=2$	$h=4$	$h=8$
CP-RFR $\alpha = 0.85$	ICP	0.86	0.85	0.85	0.85
	MIL	627.12	716.45	918.45	1214.01
E-RFR $\alpha = 0.9$	ICP	0.85	0.85	0.85	0.85
	MIL	581.63	675.8	855.98	1151.84
CP-RFR $\alpha = 0.99$	ICP	0.99	0.99	0.99	0.99
	MIL	1535.12	1896.72	2596.84	3668.75
HR $\alpha = 0.9$	ICP	0.97	0.97	0.97	0.98
	MIL	895.51	1044.98	1434.81	2144.94

in Table V. Figure 12 shows the calibration curves (also referred to as *reliability diagrams* for each of these cases). A calibration plot examines whether the confidence interval estimated for a confidence level α actually captures a fraction α of the observed test values. The area between the curve obtained for different α values and the ideal case (a perfectly calibrated model) is a numerical indicator of how miscalibrated the model can be regarded to be. Plots included in this figure reveal that both E-RFR and HR are not properly calibrated for high values of α : E-RFR is slightly overconfident in its estimated uncertainty (i.e., the computed intervals are too narrow for what they should be), whereas HR is found to be underconfident in this same region of α values (namely, its estimated intervals cover are too wide). By contrast, the calibration process performed in conformal prediction allows computing accurate interval estimates for any value of α , as exposed by the closeness of its calibration curves to the ideal case and the notably smaller miscalibration area annotated inside the plots.

V. CONCLUSION

As in many other areas related to ITS, traffic forecasting has widely embraced the irruption of data-based

modeling approaches relying on Machine Learning algorithms. Advances held over the years have achieved in a pursuit towards solutions capable of producing forecasts of ever-growing precision, exploiting efficiently relationships held within traffic data flows over space and time. Lately, performance-driven research studies are progressively steering towards the trustworthiness and actionability of traffic forecasts produced by such data-based models, in view of the narrow performance gaps attained by modern modeling choices.

In this context, this paper has aimed at bringing the attention of the community working in traffic forecasting to this matter. Among the manifold perspectives from which one can favor the trustworthiness of data-based models, we have emphasized on the need for quantifying the confidence of the data-based model associated to its predicted traffic values. Assessing the uncertainty propagated to the model's output allows a traffic manager to better design countermeasures against future congestion events in the road network, delineate better traffic light schedules, decide where to deploy new traffic sensors and collect data therefrom, or quantify whether forecasts can be predicted more confidently if the model is supplied with data supplied by new sources of information. In short, uncertainty quantification (also referred to as model's confidence) is a key for human decision making based on the output of traffic forecasting models. Besides providing this rationale, this work provides an overview on the most representative uncertainty estimation methods, as well as quantitative measures to gauge the quality of the estimated confidence intervals. This first part of the manuscript pretends to be a soft entry point to newcomers interested in confidence-aware traffic forecasting, establishing the main motivational reasons for research in the area, essential information pointers, a summary of baseline techniques and a description of evaluation measures and protocols.

To further complement this material, a comprehensive set of experiments over real traffic data collected in the city of Madrid (Spain) has been designed, comprising different sets of features, uncertainty estimation techniques and Machine Learning based forecasting models. Results stemming from this setup have lead to several lessons learned about the role of confidence/uncertainty in the actionability of traffic forecasts. In the first place, a diversity of uncertainty estimation techniques has been compared to each other, evincing that there is not a single source of uncertainty, nor is there a unique way of approaching its calculation. Nevertheless, the metrics used to assess the validity and informational efficiency of the confidence intervals estimated by such techniques has uncovered that Conformal Prediction is more stable and reliable than approaches based on the ablation of models to statistically characterize their output (e.g. ensemble models or Monte Carlo dropout), but also better than models specifically suited for the purpose (e.g., Bayesian Neural Networks). Besides, Conformal Prediction yields a transparent and traceable way of obtaining confidence intervals, and the calibration stage that lies at the core of its procedure helps maintain the size of the intervals to the minimum that guarantees the coverage of unseen test samples. This feature, combined with its model-agnostic nature, renders Conformal Prediction as a

very interesting option for estimating the uncertainty of traffic forecasts. On the other side, methods like evidential Deep Learning fail to produce reliable intervals due to its known susceptibility to the value of its parameters and the suitability of the evidential priors. Other methods yield different levels of quality in regards to their estimated confidence intervals, which depend on the characteristics present in the input data from which traffic forecasts are predicted. Our experiments have also confirmed that calibration is essential to reliably estimate the uncertainty of forecasts for a given confidence level.

Further along this line, we have verified that the particularities of the dataset at hand may affect the amount of uncertainty associated to the model's output. In the traffic dataset used for the experiments, the stability and inherent predictive potential of the traffic time series cause that other sources of data that could potentially help to deliver more precise forecasts have a low relevance and a low impact in uncertainty. In general, short-term traffic forecasting models can benefit marginally from additional sources of data, due to the acknowledged relevance of the time series variables that take part in the model. However, there are cases in which the impact becomes noticeable in the precision of traffic forecasts, whereas in other cases the addition of new data sources produces similar performance results, yet worst confidence intervals. These cases prove that the composition of the datasets used for traffic forecasting can also be driven by the examination of its consequences for the uncertainty of the model. Therefore, a principled assessment of the uncertainty of traffic forecasts throughout the modeling pipeline and the actionability that it grants should be of pivotal importance in future studies, considering it as an additional dimension in comparison benchmarks, and as a criterion to decide whether to include new sources of information (e.g. social media, traffic cameras) that may jeopardize the confidence of the model in its predictions.

Beyond the questions addressed experimentally in this research work, other applications of uncertainty estimation that were presented in Section II have not been explored to date. We envision a rich agenda, plenty of uncharted research lines, related to confidence-aware traffic forecasting. Among them, we remark the high stability of properly calibrated intervals, which can be helpful for trend change detection over traffic data streams. Indeed, a continuously flowing stream of traffic data should fall within the estimated confidence interval, since it statistically represents the *regular* behavior of traffic data in the location of interest. If real samples eventually start falling out the confidence interval, this can be symptomatic of a contextual change of the traffic behavior (due to e.g., roadworks or any exceptional circumstance like an accident), so that a closer inspection can be enforced or an update of the traffic forecasting model be triggered upon its occurrence. When dealing with spatio-temporal models such as GNNs, a question that naturally arises in the context of this manuscript is how the uncertainty present at the output of the model distributes over space and time. We foresee that the development of methods capable of attributing the prediction uncertainty to the different collection points of a road network

can be informative to design data collection strategies using temporally deployed traffic sensors. This hypothesis, together with the current momentum of spatio-temporal GNN modelling for traffic forecasting, is an interesting path to extend the insights provided in the present work. Lastly, a promising application of uncertainty estimation techniques arises from its natural connection to the active learning research area: by measuring confidence intervals over different regions of the feature space at the input of the model, one can identify where most of its uncertainty is concentrated, contributing to the identification of suitable locations in the road network and/or periods in time where more traffic data should be collected. This information can be ultimately be used for designing new traffic measurement campaigns with provisionally deployed sensors over the city, showcasing another purpose for which uncertainty estimations can effectively contribute to the actionability of traffic forecasts.

REFERENCES

- [1] I. Lana, J. Del Ser, M. Velez, and E. I. Vlahogianni, "Road traffic forecasting: Recent advances and new challenges," *IEEE Intell. Transp. Syst. Mag.*, vol. 10, no. 2, pp. 93–109, Summer 2018.
- [2] M. S. Ahmed and A. R. Cook, "Analysis of freeway traffic time-series data by using box-jenkins techniques," *Transp. Res. Rec.*, no. 722, pp. 1–9, 1979.
- [3] M. Levin and Y.-D. Tsao, "On forecasting freeway occupancies and volumes (abridgment)," *Transp. Res. Rec.*, vol. 773, pp. 47–49, Oct. 1980.
- [4] C. K. Moorthy and B. G. Ratcliffe, "Short term traffic forecasting using time series methods," *Transp. Planning Technol.*, vol. 12, no. 1, pp. 45–56, Jul. 1988.
- [5] E. I. Vlahogianni, M. G. Karlaftis, and J. C. Golias, "Spatio-temporal short-term urban traffic volume forecasting using genetically optimized modular networks," *Comput.-Aided Civil Infrastruct. Eng.*, vol. 22, no. 5, pp. 317–325, Jul. 2007.
- [6] A. M. Nagy and V. Simon, "Survey on traffic prediction in smart cities," *Pervas. Mobile Comput.*, vol. 50, pp. 148–163, Oct. 2018.
- [7] E. L. Manibardo, I. Laña, and J. Del Ser, "Transfer learning and online learning for traffic forecasting under different data availability conditions: Alternatives and pitfalls," in *Proc. IEEE 23rd Int. Conf. Intell. Transp. Syst. (ITSC)*, Sep. 2020, pp. 1–6.
- [8] A. Ermagun and D. Levinson, "Spatiotemporal traffic forecasting: Review and proposed directions," *Transp. Rev.*, vol. 38, no. 6, pp. 786–814, Nov. 2018.
- [9] W. Jiang and J. Luo, "Graph neural network for traffic forecasting: A survey," *Expert Syst. Appl.*, vol. 207, Nov. 2022, Art. no. 117921.
- [10] E. L. Manibardo, I. Laña, and J. D. Ser, "Deep learning for road traffic forecasting: Does it make a difference?" *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 7, pp. 6164–6188, Jul. 2022.
- [11] X. Yin, G. Wu, J. Wei, Y. Shen, H. Qi, and B. Yin, "Deep learning on traffic prediction: Methods, analysis, and future directions," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 6, pp. 4927–4943, Jun. 2022.
- [12] I. La na, J. J. Sanchez-Medina, E. I. Vlahogianni, and J. D. Ser, "From data to actions in intelligent transportation systems: A prescription of functional requirements for model actionability," *Sensors*, vol. 21, no. 4, p. 1121, Feb. 2021.
- [13] E. I. Vlahogianni, M. G. Karlaftis, and J. C. Golias, "Short-term traffic forecasting: Where we are and where we're going," *Transp. Res. C, Emerg. Technol.*, vol. 43, pp. 3–19, Jun. 2014.
- [14] A. Majda and N. Chen, "Model error, information barriers, state estimation and prediction in complex multiscale systems," *Entropy*, vol. 20, no. 9, p. 644, Aug. 2018.
- [15] R. Anthes and T. Rieckh, "Estimating observation and model error variances using multiple data sets," *Atmos. Meas. Techn.*, vol. 11, no. 7, pp. 4239–4260, Jul. 2018.
- [16] G. de Jong, A. Daly, M. Pieters, S. Miller, R. Plasmeijer, and F. Hofman, "Uncertainty in traffic forecasts: Literature review and new results for The Netherlands," *Transportation*, vol. 34, no. 4, pp. 375–395, Jul. 2007.
- [17] N. N. Taleb, "Errors, robustness, and the fourth quadrant," *Int. J. Forecasting*, vol. 25, no. 4, pp. 744–759, Oct. 2009.
- [18] J. Hilden, J. D. F. Habbema, and B. Bjerregaard, "The measurement of performance in probabilistic diagnosis," *Methods Inf. Med.*, vol. 17, no. 4, pp. 227–237, 1978.
- [19] H. Papadopoulos, A. Gammerman, and V. Vovk, "Confidence predictions for the diagnosis of acute abdominal pain," in *Proc. IFIP Int. Conf. Artif. Intell. Appl. Innov.* Cham, Switzerland: Springer, 2009, pp. 175–184.
- [20] F. Yang, H.-Z. Wang, H. Mi, C.-D. Lin, and W.-W. Cai, "Using random forest for reliable classification and cost-sensitive learning for medical diagnosis," *BMC Bioinf.*, vol. 10, no. S1, p. S22, Jan. 2009.
- [21] V. Balasubramanian, R. Gouripeddi, S. Panchanathan, J. Vermillion, A. Bhaskaran, and R. Siegel, "Support vector machine based conformal predictors for risk of complications following a coronary drug eluting stent procedure," in *Proc. 36th Annu. Comput. Cardiology Conf. (CinC)*, Sep. 2009, pp. 5–8.
- [22] A. Lambrou, H. Papadopoulos, and A. Gammerman, "Reliable confidence measures for medical diagnosis with evolutionary algorithms," *IEEE Trans. Inf. Technol. Biomed.*, vol. 15, no. 1, pp. 93–99, Jan. 2011.
- [23] S. Zhou, M. Mao, and J. Su, "Short-term forecasting of wind power and non-parametric confidence interval estimation," in *Zhongguo Dianji Gongcheng Xuebao(Proc. Chin. Soc. Electr. Eng.)*, vol. 31, no. 25, pp. 10–16, 2011.
- [24] V. Almeida and J. Gama, "Prediction intervals for electric load forecast: Evaluation for different profiles," in *Proc. 18th Int. Conf. Intell. Syst. Appl. Power Syst. (ISAP)*, Sep. 2015, pp. 1–6.
- [25] J. Neeven and E. Smirnov, "Conformal stacked weather forecasting," in *Conformal and Probabilistic Prediction and Applications*, vol. 91. Proceedings of Machine Learning Research, 2018, pp. 220–233. [Online]. Available: <https://proceedings.mlr.press/v91/neeven18a.html>
- [26] S. Scher and G. Messori, "Predicting weather forecast uncertainty with machine learning," *Quart. J. Roy. Meteorolog. Soc.*, vol. 144, no. 717, pp. 2830–2841, Oct. 2018.
- [27] H. Lee and B. Lee, "Confidence-aware deep learning forecasting system for daily solar irradiance," *IET Renew. Power Gener.*, vol. 13, no. 10, pp. 1681–1689, Jul. 2019.
- [28] C. Hubmann, J. Schulz, M. Becker, D. Althoff, and C. Stiller, "Automated driving in uncertain environments: Planning with interaction and uncertain maneuver prediction," *IEEE Trans. Intell. Vehicles*, vol. 3, no. 1, pp. 5–17, Mar. 2018.
- [29] S. Noh, "Decision-making framework for autonomous driving at road intersections: Safeguarding against collision, overly conservative behavior, and violation vehicles," *IEEE Trans. Ind. Electron.*, vol. 66, no. 4, pp. 3275–3286, Apr. 2019.
- [30] L. Zhang, W. Ding, J. Chen, and S. Shen, "Efficient uncertainty-aware decision-making for automated driving using guided branching," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2020, pp. 3291–3297.
- [31] X. Huang, S. G. McGill, B. C. Williams, L. Fletcher, and G. Rosman, "Uncertainty-aware driver trajectory prediction at urban intersections," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2019, pp. 9718–9724.
- [32] C. Choi, S. Malla, A. Patil, and J. Hee Choi, "DROGON: A trajectory prediction model based on intention-conditioned behavior reasoning," 2019, *arXiv:1908.00024*.
- [33] J. Li, H. Ma, W. Zhan, and M. Tomizuka, "Coordination and trajectory prediction for vehicle interactions via Bayesian generative modeling," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Jun. 2019, pp. 2496–2503.
- [34] E. J. Casado Magaña, "Trajectory prediction uncertainty modelling for air traffic management," Ph.D. dissertation, Division Aerosp. Sci., School Eng., Univ. Glasgow, Glasgow, Scotland, 2016.
- [35] P. Scala, M. M. Mota, J. Ma, and D. Delahaye, "Tackling uncertainty for the development of efficient decision support system in air traffic management," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 8, pp. 3233–3246, Aug. 2020.
- [36] J. Chen, L. Chen, and D. Sun, "Air traffic flow management under uncertainty using chance-constrained optimization," *Transp. Res. B, Methodol.*, vol. 102, pp. 124–141, Aug. 2017.
- [37] P. Parthasarathi and D. Levinson, "Post-construction evaluation of traffic forecast accuracy," *Transp. Policy*, vol. 17, no. 6, pp. 428–443, Nov. 2010.
- [38] M. S. Nicolaisen and P. A. Driscoll, "Ex-Post Evaluations of demand forecast accuracy: A literature review," *Transp. Rev.*, vol. 34, no. 4, pp. 540–557, Jul. 2014.

- [39] C. Yang, A. Chen, X. Xu, and S. Wong, "Sensitivity-based uncertainty analysis of a combined travel demand model," *Transp. Res. B, Methodol.*, vol. 57, pp. 225–244, Nov. 2013.
- [40] A. Matas, J.-L. Raymond, and A. Ruiz, "Traffic forecasts under uncertainty and capacity constraints," *Transportation*, vol. 39, no. 1, pp. 1–17, Jan. 2012.
- [41] M. B. Hugosson, "Quantifying uncertainties in a national forecasting model," *Transp. Res. A, Policy Pract.*, vol. 39, no. 6, pp. 531–547, Jul. 2005.
- [42] S. Rasouli and H. J. P. Timmermans, "Using ensembles of decision trees to predict transport mode choice decisions: Effects on predictive success and uncertainty estimates," *Eur. J. Transp. Infrastruct. Res.*, vol. 14, no. 4, pp. 412–424, 2014.
- [43] M. Welde and J. Odeck, "Do planners get it right? The accuracy of travel demand forecasting in Norway," *Eur. J. Transp. Infrastruct. Res.*, vol. 11, no. 1, pp. 1–16, 2011.
- [44] I. Laña, E. Villar-Rodriguez, U. Etxegarai, I. Oregi, and J. D. Ser, "A question of trust: Statistical characterization of long-term traffic estimations for their improved actionability," in *Proc. IEEE Intell. Transp. Syst. Conf. (ITSC)*, Oct. 2019, pp. 1922–1928.
- [45] R. C. Smith, *Uncertainty Quantification: Theory, Implementation, and Applications*, vol. 12. Philadelphia, PA, USA: SIAM, 2013.
- [46] T. J. Sullivan, *Introduction to Uncertainty Quantification*, vol. 63. Cham, Switzerland: Springer, 2015.
- [47] J. Guo and B. M. Williams, "Real-time short-term traffic speed level forecasting and uncertainty quantification using layered Kalman filters," *Transp. Res. Rec., J. Transp. Res. Board*, vol. 2175, no. 1, pp. 28–37, Jan. 2010.
- [48] J. Guo, W. Huang, and B. M. Williams, "Adaptive Kalman filter approach for stochastic short-term traffic flow rate prediction and uncertainty quantification," *Transp. Res. C, Emerg. Technol.*, vol. 43, pp. 50–64, Jun. 2014.
- [49] A. B. Arrieta et al., "Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," *Inf. Fusion*, vol. 58, pp. 82–115, Jun. 2020.
- [50] K. Beven, "Facets of uncertainty: Epistemic uncertainty, non-stationarity, likelihood, hypothesis testing, and communication," *Hydrolog. Sci. J.*, vol. 61, no. 9, pp. 1652–1665, Jul. 2016.
- [51] A. D. Kiureghian and O. Ditlevsen, "Aleatory or epistemic? Does it matter?" *Structural Saf.*, vol. 31, no. 2, pp. 105–112, Mar. 2009.
- [52] D. V. Lindley, "The philosophy of statistics," *J. Roy. Stat. Soc., Ser. D, Statistician*, vol. 49, no. 3, pp. 293–337, 2000.
- [53] A. Shapcev, K. Gubaev, E. Tsymbalov, and E. Podryabinkin, "Active learning and uncertainty estimation," in *Machine Learning Meets Quantum Physics* (Lecture Notes in Physics), vol. 968, K. Schütt, S. Chmiela, O. von Lilienfeld, A. Tkatchenko, K. Tsuda, and K. R. Müller, Eds. Cham, Switzerland: Springer, 2020, doi: 10.1007/978-3-030-40245-7_15.
- [54] J. Zhu, H. Wang, B. K. Tsou, and M. Ma, "Active learning with sampling by uncertainty and density for data annotations," *IEEE Trans. Audio, Speech, Language Process.*, vol. 18, no. 6, pp. 1323–1331, Aug. 2010.
- [55] Y. Yang and M. Loog, "Active learning using uncertainty information," in *Proc. 23rd Int. Conf. Pattern Recognit. (ICPR)*, Dec. 2016, pp. 2646–2651.
- [56] V.-L. Nguyen, M. H. Shaker, and E. Hüllermeier, "How to measure uncertainty in uncertainty sampling for active learning," *Mach. Learn.*, vol. 111, no. 1, pp. 89–122, Jan. 2022.
- [57] E. L. Manibardo, I. Laña, E. Villar-Rodriguez, and J. D. Ser, "A graph-based methodology for the sensorless estimation of road traffic profiles," *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 8, pp. 8701–8715, Aug. 2023.
- [58] L. Baier, T. Schlör, J. Schöffner, and N. Kühl, "Detecting concept drift with neural network model uncertainty," 2021, *arXiv:2107.01873*.
- [59] S. Ghosh et al., "Uncertainty quantification 360: A holistic toolkit for quantifying and communicating the uncertainty of AI," 2021, *arXiv:2106.01410*. [Online]. Available: <https://arxiv.org/abs/2106.01410>
- [60] E. Hüllermeier and W. Waegeman, "Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods," *Mach. Learn.*, vol. 110, no. 3, pp. 457–506, 2021.
- [61] R. Senge et al., "Reliable classification: Learning classifiers that distinguish aleatoric and epistemic uncertainty," *Inf. Sci.*, vol. 255, pp. 16–29, Jan. 2014.
- [62] D. Wu et al., "Quantifying uncertainty in deep spatiotemporal forecasting," in *Proc. 27th ACM SIGKDD Conf. Knowl. Discovery Data Mining*, Aug. 2021, pp. 1841–1851.
- [63] M. Abdar et al., "A review of uncertainty quantification in deep learning: Techniques, applications and challenges," *Inf. Fusion*, vol. 76, pp. 243–297, Dec. 2021.
- [64] A. Gammerman, V. Vovk, and V. Vapnik, "Learning by transduction," 2013, *arXiv:1301.7375*.
- [65] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic Learning in a Random World*. Cham, Switzerland: Springer, 2005.
- [66] V. Balasubramanian, S.-S. Ho, and V. Vovk, *Conformal Prediction for Reliable Machine Learning: Theory, Adaptations and Applications*. San Francisco, CA, USA: Newnes, 2014.
- [67] U. Johansson, H. Boström, T. Löfström, and H. Linusson, "Regression conformal prediction with random forests," *Mach. Learn.*, vol. 97, nos. 1–2, pp. 155–176, Oct. 2014.
- [68] T. G. Dietterich et al., "Ensemble learning," in *The Handbook Brain Theory Neural Networks*, vol. 2. Cambridge, MA, USA: MIT Press, 2002, pp. 110–125.
- [69] S. González, S. García, J. Del Ser, L. Rokach, and F. Herrera, "A practical tutorial on bagging and boosting based ensembles for machine learning: Algorithms, software tools, performance study, practical perspectives and opportunities," *Inf. Fusion*, vol. 64, pp. 205–237, Dec. 2020.
- [70] L. Mentch and G. Hooker, "Quantifying uncertainty in random forests via confidence intervals and hypothesis tests," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 841–881, 2016.
- [71] D. A. Freedman, "Bootstrapping regression models," *Ann. Statist.*, vol. 9, no. 6, pp. 1218–1228, Nov. 1981.
- [72] H. Sun, C. Zhang, and B. Ran, "Interval prediction for traffic time series using local linear predictor," in *Proc. 7th Int. IEEE Conf. Intell. Transp. Syst.*, Jul. 2004, pp. 410–415.
- [73] B. Naik, L. R. Rilett, J. Appiah, and L. F. Walubita, "Resampling methods for estimating travel time uncertainty: Application of the gap bootstrap," *Transp. Res. Rec., J. Transp. Res. Board*, vol. 2672, no. 42, pp. 137–147, Dec. 2018.
- [74] S. Kumar and A. N. Srivastava, "Bootstrap prediction intervals in non-parametric regression with applications to anomaly detection," in *Proc. 18th ACM SIGKDD Conf. Knowl. Discovery Data Mining*, 2012, pp. 1–9.
- [75] R. Koenker and K. Hallock, "Quantile regression," *J. Econ. Perspect.*, vol. 15, no. 4, pp. 143–156, 2001.
- [76] O. Rahmati et al., "Predicting uncertainty of machine learning models for modelling nitrate pollution of groundwater using quantile regression and UNEEC methods," *Sci. Total Environ.*, vol. 688, pp. 855–866, Oct. 2019.
- [77] K. Vaysse and P. Lagacherie, "Using quantile regression forest to estimate uncertainty of digital soil mapping products," *Geoderma*, vol. 291, pp. 55–64, Apr. 2017.
- [78] R. Wen, K. Torkkola, B. Narayanaswamy, and D. Madeka, "A multi-horizon quantile recurrent forecaster," 2017, *arXiv:1711.11053*.
- [79] C. Fan et al., "Multi-horizon time series forecasting with temporal attention learning," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2019, pp. 2527–2535.
- [80] H. Wang and D.-Y. Yeung, "Towards Bayesian deep learning: A framework and some existing methods," *IEEE Trans. Knowl. Data Eng.*, vol. 28, no. 12, pp. 3395–3408, Dec. 2016.
- [81] A. A. Abdullah, M. M. Hassan, and Y. T. Mustafa, "A review on Bayesian deep learning in healthcare: Applications and challenges," *IEEE Access*, vol. 10, pp. 36538–36562, 2022.
- [82] S. Hernández and J. L. López, "Uncertainty quantification for plant disease detection using Bayesian deep learning," *Appl. Soft Comput.*, vol. 96, Nov. 2020, Art. no. 106597.
- [83] H. Lee and B.-T. Lee, "Bayesian deep learning-based confidence-aware solar irradiance forecasting system," in *Proc. Int. Conf. Inf. Commun. Technol. Converg. (ICTC)*, Oct. 2018, pp. 1233–1238.
- [84] B. Wang et al., "Deep uncertainty quantification: A machine learning approach for weather forecasting," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discovery & Data Mining*, 2019, pp. 2087–2095.
- [85] Y. Wu and J. J. Q. Yu, "A Bayesian learning network for traffic speed forecasting with uncertainty quantification," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2021, pp. 1–7.
- [86] F. Zhou, Q. Yang, T. Zhong, D. Chen, and N. Zhang, "Variational graph neural networks for road traffic prediction in intelligent transportation systems," *IEEE Trans. Ind. Informat.*, vol. 17, no. 4, pp. 2802–2812, Apr. 2021.

- [87] D. Salinas, V. Flunkert, J. Gasthaus, and T. Januschowski, "DeepAR: Probabilistic forecasting with autoregressive recurrent networks," *Int. J. Forecast.*, vol. 36, no. 3, pp. 1181–1191, 2020.
- [88] A. Amini, W. Schwarting, A. Soleimany, and D. Rus, "Deep evidential regression," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 14927–14937.
- [89] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1050–1059.
- [90] R. L. Winkler et al., "Scoring rules and the evaluation of probabilities," *Test*, vol. 5, no. 1, pp. 1–60, Jun. 1996.
- [91] B. Efron and T. Hastie, *Computer Age Statistical Inference, Student Edition: Algorithms, Evidence, and Data Science*. Cambridge, U.K.: Cambridge Univ. Press, 2021.
- [92] A. Kendall and Y. Gal, "What uncertainties do we need in Bayesian deep learning for computer vision?" in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 1–11.
- [93] J. Fu, W. Zhou, and Z. Chen, "Bayesian graph convolutional network for traffic prediction," 2021, [arXiv:2104.00488](https://arxiv.org/abs/2104.00488).
- [94] J. Xia, S. Wang, X. Wang, M. Xia, K. Xie, and J. Cao, "Multi-view Bayesian spatio-temporal graph neural networks for reliable traffic flow prediction," *Int. J. Mach. Learn. Cybern.*, vol. 15, pp. 1–14, Oct. 2022.
- [95] S. Luan, R. Ke, Z. Huang, and X. Ma, "Traffic congestion propagation inference using dynamic Bayesian graph convolution network," *Transp. Res. C, Emerg. Technol.*, vol. 135, Feb. 2022, Art. no. 103526.
- [96] L. Du, R. Gao, P. N. Suganthan, and D. Z. W. Wang, "Graph ensemble deep random vector functional link network for traffic forecasting," *Appl. Soft Comput.*, vol. 131, Dec. 2022, Art. no. 109809.
- [97] S.-Y. Han, Q. Zhao, Q.-W. Sun, J. Zhou, and Y.-H. Chen, "EnGS-DGR: Traffic flow forecasting with indefinite forecasting interval by ensemble GCN, Seq2Seq, and dynamic graph reconfiguration," *Appl. Sci.*, vol. 12, no. 6, p. 2890, Mar. 2022.
- [98] T. Mallick, P. Balaprakash, and J. Macfarlane, "Deep-ensemble-based uncertainty quantification in spatiotemporal graph neural networks for traffic forecasting," 2022, [arXiv:2204.01618](https://arxiv.org/abs/2204.01618).
- [99] P. Shang, X. Liu, C. Yu, G. Yan, Q. Xiang, and X. Mi, "A new ensemble deep graph reinforcement learning network for spatio-temporal traffic volume forecasting in a freeway network," *Digit. Signal Process.*, vol. 123, Apr. 2022, Art. no. 103419.
- [100] Z. Feng, K. Qi, B. Shi, H. Mei, Q. Zheng, and H. Wei, "Deep evidential learning in diffusion convolutional recurrent neural network," *Electron. Res. Arch.*, vol. 31, no. 4, pp. 2252–2264, 2023.
- [101] F. Rodrigues and F. C. Pereira, "Heteroscedastic Gaussian processes for uncertainty modeling in large-scale crowdsourced traffic data," *Transp. Res. C, Emerg. Technol.*, vol. 95, pp. 636–651, Oct. 2018.
- [102] K. Tran, W. Neiswanger, J. Yoon, Q. Zhang, E. Xing, and Z. W. Ulissi, "Methods for comparing uncertainty quantifications for material property predictions," *Mach. Learn., Sci. Technol.*, vol. 1, no. 2, Jun. 2020, Art. no. 025006.
- [103] N. C. Petersen, A. Parslov, and F. Rodrigues, "Short-term bus travel time prediction for transfer synchronization with intelligent uncertainty handling," 2021, [arXiv:2104.06819](https://arxiv.org/abs/2104.06819).
- [104] D. L. Shrestha and D. P. Solomatine, "Machine learning approaches for estimation of prediction interval for the model output," *Neural Netw.*, vol. 19, no. 2, pp. 225–235, Mar. 2006.
- [105] T. Gneiting, F. Balabdaoui, and A. E. Raftery, "Probabilistic forecasts, calibration and sharpness," *J. Roy. Stat. Soc. Ser. B, Stat. Methodol.*, vol. 69, no. 2, pp. 243–268, Apr. 2007.
- [106] T. Gneiting and A. E. Raftery, "Strictly proper scoring rules, prediction, and estimation," *J. Amer. Stat. Assoc.*, vol. 102, no. 477, pp. 359–378, Mar. 2007.
- [107] A. Carvalho, "An overview of applications of proper scoring rules," *Decis. Anal.*, vol. 13, no. 4, pp. 223–242, Dec. 2016.
- [108] R. E. Schapire, "Explaining AdaBoost," in *Empirical Inference*, N. V. Vladimir, Ed. Berlin, Germany: Springer, 2013, pp. 37–52.
- [109] I. Laña, J. Del Ser, A. Padró, M. Vélez, and C. Casanova-Mateo, "The role of local urban traffic and meteorological conditions in air pollution: A data-based case study in madrid, Spain," *Atmos. Environ.*, vol. 145, pp. 424–438, Nov. 2016.
- [110] E. L. Manibardo, I. Laña, and J. Del Ser, "What to sense when there is no sensor: Ex-novo traffic flow estimation for non-sensed roads," in *Proc. IEEE 25th Int. Conf. Intell. Transp. Syst. (ITSC)*, Oct. 2022, pp. 1668–1675.
- [111] H. V. Gupta, K. J. Beven, and T. Wagener, "Model calibration and uncertainty estimation," in *Encyclopedia of Hydrological Sciences*. Hoboken, NJ, USA: Wiley, 2006, ch. 131, doi: [10.1002/0470848944.hsa138](https://doi.org/10.1002/0470848944.hsa138).



Ibai Laña (Member, IEEE) received the bachelor's degree in computer science from the University of Deusto, Spain, the M.Sc. degree in artificial intelligence from Universidad Nacional de Educación a Distancia (UNED), Spain, and the Ph.D. degree (cum laude) in mobile network information and communication technologies from the University of the Basque Country, Spain. He is currently a Principal Investigator at TECNALIA and a Professor with the Department of Computer, Electronic and Communication Technologies, Deusto University. He has

been a Visiting Scholar with the Knowledge Engineering and Discovery Research Institute (KEDRI), Auckland University of Technology (AUT). He is the author of more than 50 publications and one patent. His research interests include intelligent transportation systems (ITS), machine learning, traffic data analysis, and data science.



Ignacio (Iñaki) Olabarrieta received the bachelor's degree in science from the University of the Basque Country, Spain, and the M.Sc. and Ph.D. degrees in physics from The University of British Columbia, Canada. He was a Post-Doctoral Researcher with Hearne Institute of Theoretical Physics, Louisiana State University (LSU), Baton Rouge, LA, USA, where he focused on modeling relativistic systems by computational means. Since 2006, he has been a Researcher with TECNALIA. He is also a principal investigator in the realm of artificial intelligence with

a special focus on its application to mobility. He is the author of many research articles and two patents. He has taken part in numerous European-funded projects.



Javier Del Ser (Senior Member, IEEE) received the first Ph.D. degree (cum laude) in control engineering and industrial electronics from the University of Navarra, Spain, and the second Ph.D. degree (cum laude) in information and communication technologies from the University of Alcalá de Henares, Spain, in 2013. He is currently a Professor with the Department of Communications Engineering, University of the Basque Country (EHU/UPV), and a Research Professor and the Chief Scientist of Applied Artificial Intelligence with TECNALIA.

He is also a Distinguished Professor with the University of Granada, Spain. He has published approximately 500 scientific-technical publications to date in these fields, edited four books, supervised 16 Ph.D. theses, and participated/directed more than 60 projects and contracts. His research interests include artificial intelligence, machine learning, and deep learning applied to practical modeling and optimization tasks for problems in different sectors, with an emphasis on data-based modeling in mobility, logistics, and transportation. He has been included in the list of the top 2% most influential AI researchers worldwide by Stanford University (2021, 2022, and 2023). He has also been part of the team that developed the R&D&i strategy in artificial intelligence for the Government of Spain in 2019. He received the Extraordinary Ph.D. Award for the Ph.D. degree. He is an Associate Editor of IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS.