

Poster Abstract: Predictive Monitoring with Uncertainty for Deep Learning Enabled Smart Cities

Meiyi Ma
University of Virginia
meiyi@virginia.edu

Ezio Bartocci
TU Wien
ezio.bartocci@tuwien.ac.at

John Stankovic, Lu Feng
University of Virginia
stankovic,lu.feng@virginia.edu

ABSTRACT

In order to prevent safety violations, predictive monitoring with uncertainty is crucial for deep learning-enabled services in smart cities. We develop a novel predictive monitoring system for smart city applications, which consists of an RNN-based predictor with uncertainty estimation and a new specification language, named Signal Temporal Logic with Uncertainty. The solution first predicts a sequence of distributions representing city's future states with uncertainty estimation and then checks the predicted results against STL-U specified safety and performance requirements. The system supports decision making by providing a quantitative satisfaction degree with confidence guarantees. We receive promising results from evaluations on two large-scale city datasets, and on a case study on real-time predictive monitoring in a simulated smart city.

CCS CONCEPTS

• **Computer systems organization** → **Embedded and cyber-physical systems**; • **Computing methodologies** → **Bayesian network models**.

KEYWORDS

Deep Learning, Uncertainty, Predictive Monitoring, Smart Cities

ACM Reference Format:

Meiyi Ma, Ezio Bartocci, and John Stankovic, Lu Feng. 2020. Poster Abstract: Predictive Monitoring with Uncertainty for Deep Learning Enabled Smart Cities. In *The 18th ACM Conference on Embedded Networked Sensor Systems (SenSys '20)*, November 16–19, 2020, Virtual Event, Japan. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3384419.3430445>

1 INTRODUCTION

Smart services are increasingly embedded in modern cities aiming to enhance various aspects of citizens' lives, including safety, well-being, and quality of life [2]. Smart services infer city states from various sensors deployed in the city, and take actions on actuators. Meanwhile, City Operations Centers (e.g., IBMs Intelligent Operations Center and Microsoft CityNext) deploy control platforms to monitor city states and services against city safety and performance requirements to make decisions. If a requirement violation is detected by the monitor, the city operators can take actions to change

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SenSys '20, November 16–19, 2020, Virtual Event, Japan

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7590-0/20/11...\$15.00

<https://doi.org/10.1145/3384419.3430445>

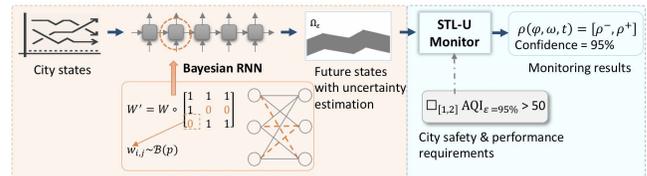


Figure 1: System Overview

the states, such as controlling traffic signals, rejecting unsafe actions, sending alarms to police, etc. To prevent unsafe situations ahead of time, more and more services and control centers predict city's future states with deep learning prediction models, e.g., fire risk prediction, conflict detection and resolution [3], etc.

However, for deep prediction models, with the same set of inputs, the output will always be the same. It is difficult to assess if a model is making sensible predictions just with a single value and making a decision on that single value can be unreliable. In reality, data and models are not perfect and contain uncertainties (e.g., sensing and environment noise, unexpected events, accidents) that are not reflected by the deterministic deep neural networks. Uncertainty will affect the prediction and monitoring results, and therefore influence the safety and decision making in the cities. Meanwhile, while there has been a great effort in developing monitoring techniques for smart cities such as using Signal Temporal Logic and its extensions, existing works mostly focus on monitoring a single multi-variable signal, and cannot be directly used to monitor the signals output from deep learning with uncertainty.

In this abstract, we briefly describe a novel predictive monitoring system considering uncertainties for deep learning enabled smart cities in order to assure services' safety and reliability. The system includes two major technical components (as shown in Figure 1), i.e., deep learning prediction with uncertainty, and STL-U monitoring with uncertainty. First, we cast deep learning as Bayesian models, which return a predicted sequence of city future states with uncertainty estimation (defined as flowpipes). Second, we develop a new specification language *Signal Temporal Logic with Uncertainty* (STL-U) for the monitoring of flowpipes output by predictive Bayesian deep learning models. We evaluate the system on large-scale real city datasets and perform a case study on real-time predictive monitoring for simulated New York City.

2 SYSTEM DESIGN

Our system is envisioned to be deployed in the control center of a smart city to predict the city's future states with uncertainty estimation, and then verify if the predicted results satisfy city safety and performance requirements. The monitoring results can support city decision making.

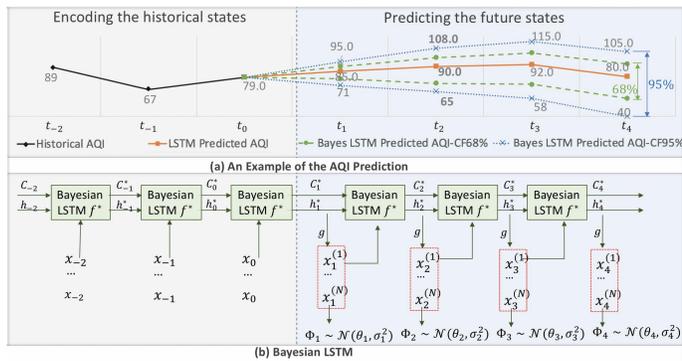


Figure 2: AQI prediction with Bayesian LSTM)

2.1 Deep Learning Prediction with Uncertainty

At runtime, the system first takes city’s current states as input, and predicts city’s future states using Bayesian RNN models [1], which considers data and model uncertainty. The predicted results are presented as a time-series flowpipe containing the city’s future N -hour states. At each time point, instead of a single value, it shows a range of the potential values under a given confidence level. To be more specific, we cast the LSTM predictor as a probabilistic model with $w = W, w \sim p(w)$ as a random variable. Bayesian LSTM predictor f^* is shown in Figure 2(b). The Bayesian LSTM formula is updated as, $x_{t+1}^* = g(h_t^*)$, and $h_t^*, c_t^* = f^*(x_t, h_{t-1}^*, c_{t-1}^*; w)$. To estimate the posterior probability distribution, we apply Monte Carlo (MC) estimate by repeating the prediction for N times. At each iteration, with the exact same input, we obtain a different set of outputs (i.e., a time series trace) $\{x_1^{(j)}, x_2^{(j)}, x_3^{(j)}, x_4^{(j)}\}$. In total, we have a set of traces ω containing N different traces. At a single time unit $t (t > 0)$, we obtain N observations of $x_t^{(j)}$ with $i = \{1, \dots, N\}$. We estimate the expectation θ_t and variance σ_t from the samples. Therefore, at time t , we obtain the estimated results from the Bayesian LSTM as a distribution, i.e. $\Phi_t \sim \mathcal{N}(\theta_t, \sigma_t^2)$.

2.2 STL-U Monitoring with Uncertainty

We develop a new specification language *Signal Temporal Logic with Uncertainty* (STL-U) for monitoring the prediction results (such as Figure 2(a)) from deep learning models such as Bayesian LSTMs against the system requirements. We first formally define new types of signals called *flowpipes* that characterize the prediction results of deep learning with uncertainty. A flowpipe signal Ω is defined over a finite discrete time domain \mathbb{T} such that $\Omega[t] = \Phi_t$ at any time $t \in \mathbb{T}$, where Φ_t is a Gaussian distribution $\mathcal{N}(\theta_t, \sigma_t^2)$. Given a confidence level $\varepsilon \in [0, 1] \subseteq \mathbb{R}$, the flowpipe at time t is bounded by a confidence interval $[\Phi_t^-(\varepsilon), \Phi_t^+(\varepsilon)]$. Then, we define the syntax and semantics of STL-U based on flowpipe signals. We denote by $\omega : \mathbb{T} \rightarrow \{\Omega\}^n$ a multi-dimensional flowpipe signal, where $\mathbb{T} = [0, d] \subseteq \mathbb{R}$ represents for a finite discrete time domain and $n = |X|$ for a finite set of (independent) *real* variables X . Each real variable $x \in X$ has a corresponding flowpipe Ω_x , whose value follows a Gaussian distribution $\Omega_x[t]$ at time t . The syntax of STL-U is defined by the grammar,

$$\varphi := \mu_x(\varepsilon) \mid \neg \varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \mathcal{U}_I \varphi_2,$$

STL-U can be used to specify system critical requirements with satisfaction confidence levels, for example, “with 90% confidence

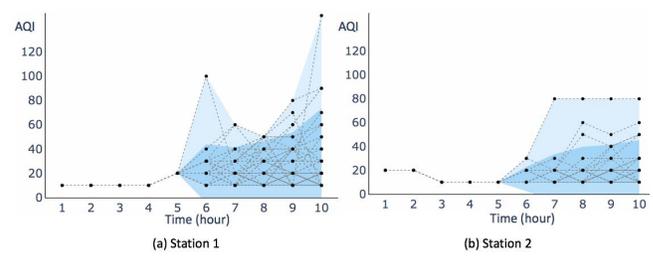


Figure 3: Uncertainty in AQI data (the dark blue shadow covers the range of 95% percentile.)

level, the predicted AQI in the next 10 hours should always be below 100”. We will define *strong* and *weak* satisfaction relations for the Boolean semantics of STL-U to indicate whether *all* or *partial* values within a flowpipe confidence interval range satisfies a requirement, respectively. In addition, we will also develop new methods for calculating the confidence level that guarantees a city requirement is satisfied by the given prediction sequence.

3 EVALUATION

We evaluate the system on large-scale real city datasets (traffic volume in NYC and AQI in Beijing) and perform case studies on real-time predictive monitoring for simulated New York City. From the preliminary results, we observe that,

(1) By analyzing the datasets (as shown in Figure 3), we found that significant uncertainty exists in the city data; the levels of uncertainty vary by datasets and locations; and pre-knowledge could also lead to different levels of uncertainty.

(2) On the evaluation of real city data, our system is effective in monitoring the flowpipes resulting from deep learning models with uncertainty. It is efficient to monitor a large scale of flowpipes resulting from deep learning models with uncertainty, e.g., on average, it only takes 417 seconds to monitor 130,000 flowpipes with an 8-time unites prediction against 390 requirements, which cannot be monitored by STL or its variants.

(3) On the real-time simulation, STL-U predictive monitor reduces the number of false violation detection, and improves the safety and performance of a smart city comparing to the monitor without considering uncertainty. For example, the air quality index is reduced by 23.7%, and emergency waiting time is reduced by 28.3% comparing to the monitor without considering uncertainty.

ACKNOWLEDGMENTS

This research was partially supported by NSF grants CCF-1942836 and CNS-1739333, and the Commonwealth Cyber Initiative, an investment from the Commonwealth of Virginia in the advancement of cyber R&D, innovation, and workforce development.

REFERENCES

- [1] Yarin Gal. 2016. *Uncertainty in deep learning*. Ph.D. Dissertation. PhD thesis, University of Cambridge.
- [2] Meiyi Ma, Sarah M Preum, Mohsin Ahmed, William Tärneberg, Abdeltawab Hendawi, and John Stankovic. 2019. Data sets, modeling, and decision making in smart cities: A survey. *ACM Transactions on Cyber-Physical Systems* 4, 2 (2019), 1–28.
- [3] Meiyi Ma, John A Stankovic, and Lu Feng. 2018. Cityresolver: a decision support system for conflict resolution in smart cities. In *Proceedings of the 9th ACM/IEEE International Conference on Cyber-Physical Systems*. IEEE Press, 55–64.