

Managing Large Dataset Gaps in Urban Air Quality Prediction: DCU-Insight-AQ at MediaEval 2022

Dinh Viet Cuong^{1,2}, Phuc H. Le-Khac², Adam Stapleton², Elke Eichlemann³, Mark Roantree^{1,2} and Alan F. Smeaton^{1,2,*}

¹*Insight Centre for Data Analytics, Dublin City University, Ireland*

²*School of Computing, Dublin City University, Ireland*

³*School of Biology and Environmental Science, University College Dublin, Ireland*

Abstract

Calculating an Air Quality Index (AQI) typically uses data streams from air quality sensors deployed at fixed locations and the calculation is a real time process. If one or a number of sensors are broken or offline, then the real time AQI value cannot be computed. Estimating AQI values for some point in the future is a predictive process and uses historical AQI values to train and build models. In this work we focus on gap filling in air quality data where the task is to predict the AQI at 1, 5 and 7 days into the future. The scenario is where one or a number of air, weather and traffic sensors are offline and explores prediction accuracy under such situations. The work is part of the MediaEval'2022 Urban Air: Urban Life and Air Pollution task submitted by the DCU-Insight-AQ team and uses multimodal and crossmodal data consisting of AQI, weather and CCTV traffic images for air pollution prediction.

1. Introduction

The Urban Life and Air Pollution task at MediaEval 2022 required participants to predict the air quality index (AQI) value at +1, +5 and +7 days using an archive of air quality, weather and images from 16 CCTV cameras, one image taken every 60 seconds [1]. Participating groups were required to download the data from online sources for local processing. Gaps in air quality datasets are common with the problem exacerbated for data gathered in poorer or developing countries [2, 3, 4, 5, 6]. In this paper we describe how we addressed the very large gaps in data that we encountered in the data we downloaded.

2. Methodology

The biggest issue in this research challenge was the significant gaps in the training data, highlighted in Figure 1 for the data we downloaded, and regarded as a common issue with climate datasets. Because participants downloaded data independently, and because data servers had different periods of downtime, it is likely that participants have different, but overlapping training data and perhaps other participants managed to download more data than we did. Even so, the amount of data we downloaded allows us to focus on the challenge of data gaps. This is

MediaEval'22: Multimedia Evaluation Workshop, January 13–15, 2023, Bergen, Norway and Online

*Corresponding author.

✉ dinh.cuong2@mail.dcu.ie (D. V. Cuong); khac.le2@mail.dcu.ie (P. H. Le-Khac); adam.stapleton9@mail.dcu.ie (A. Stapleton); elke.eichlemann@ucd.ie (E. Eichlemann); mark.roantree@dcu.ie (M. Roantree); alan.smeaton@dcu.ie (A. F. Smeaton)

🆔 0000-0002-0504-5844 (P. H. Le-Khac); 0000-0003-1233-211X (A. Stapleton); 0000-0001-9516-7951 (E. Eichlemann); 0000-0002-1329-2570 (M. Roantree); 0000-0003-1028-8389 (A. F. Smeaton)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

directly addressed in our research methodology by first identifying the sensitivity of the data gaps and adopting a counter-measure to eliminate gap data. Our method comprises 4 steps:

- Step 1. Data Analysis. This performs a statistical summary of datasets including the computation of spatial data related to the locations of air quality stations and cameras.
- Step 2. Gap Filling. Elimination or maximising the reduction in the gaps in air quality data.
- Step 3. Processing CCTV camera images. This step transforms each image into a set of features that can be combined with the air quality feature set.
- Step 4. Model Building. This step builds an experimental platform using different machine learning model configurations together with different feature sets to identify the best performing model/feature set combination.

In the remainder of this section, we described the first 3 steps in detail and in Section 3, we describe our approach to model building for the air quality prediction task.

2.1. Data Analysis

The downloaded air quality data are collected at 10 monitoring stations in Dalat City, Vietnam from March 2020 to 7th Nov 2022. The data includes air pollutant concentration for NO₂(ppm), CO, SO₂, O₃, PM1.0, PM2.5, PM10 as well as environmental measures namely temperature, humidity, UV, rainfall. In addition, traffic data, in form of images, was recorded every minute from 16 CCTV cameras across Dalat City. Figure 1 shows the availability of the dataset as downloaded by our group. This shows huge gaps in data availability. In our model building we use the first 80% of available data for training machine learning models and the remaining 20% for validation.

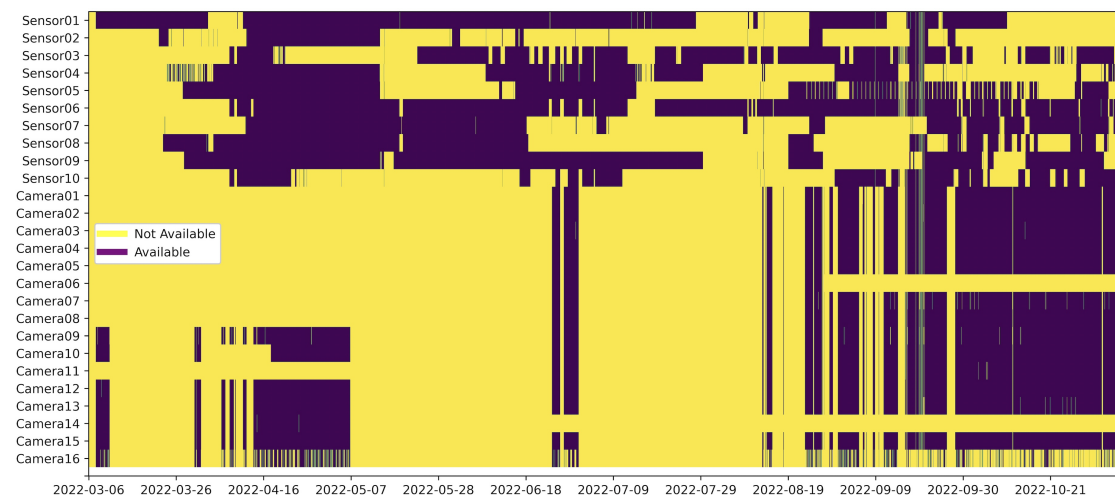


Figure 1: Missing data (shown in yellow) from across all 10 air quality measurement stations and 16 CCTV cameras for an 8-month period, early March to early November 2022.

2.2. Image Processing

For the CCTV data, we downloaded a total of 398,412 images from across all 16 cameras, which took approximately 215GB of storage. If all data had been available and downloaded, there

would be approximately 16 cameras x 8 months x 30 days x 24 hours x 60 minutes = 5,529,600 images so our download represents 7.2% of the theoretical maximum.

We re-sized each image to 640x640 and processed each using a medium-sized YOLOv6 object-detector [7] pre-trained on the COCO dataset [8]. This performs well with a balanced trade-off between speed and accuracy. From an output with more than 80 object categories, we used the average of 4 vehicle types as a proxy for traffic volume with an average detection per image of 2.58 (cars), 3.90 (motorcycles), 0.16 (bus) and 0.25 (trucks). These values are used directly as features for our predictive model without further post-processing.

2.3. Gap Filling

Our approach to gap filling used 3 different feature sets.

- Rolling Window (FS1). We generate training data using rolling windows sliding hourly over the data. Following some experimentation, we determine the historical data length to input to the model (size of window) to be 2 days. To address missing data, we remove windows with big gaps and fill windows with small gaps of less than 30% or have data from more than 4 stations.
- Rolling Window with Gap Filling (FS2). Spatial interpolation using a tree-based gradient boosting model [9] was used to infer missing values. A separate model to predict each of the air pollutants at each station was trained using all available data from other stations at that timestamp. The LightGBM model carries the same sparsity-aware learning methods as XGBoost [10] while improving on efficiency. Such methods for handling sparse arrays allow a model to learn from incomplete data.
- Rolling Window with Gap Filling and Image/Traffic Features (FS3). This feature set uses the output from the count of the numbers of cars, trucks, motorcycles and buses detected in the CCTV images from the 16 cameras.

We convert all pollutant concentration values into AQI values according to the formula described in [11]. Air quality, environmental and CCTV features are aggregated into hourly mean values which are normalised to a mean of 0.0 and standard deviation of 1.0 before input into a selection of different machine learning models.

3. Experiments

3.1. Experimental Approach

We used baseline machine learning models including multi-layer perceptrons, long short-term memory (LSTM) [12] and LSTM-GNN [13] where the prediction is treated as timeseries forecasting. We also used a spatio-temporal graph neural network [13] and our own neural network architecture labelled TemAtt and SpaTemAtt, designed to capture more complex temporal and spatio-temporal patterns.

3.2. Results

Table 1 presents results after our own validation using bold font to highlight the best performing model with the second best results underlined. The first 2 columns show the features and models used where runs marked 1-5 represent our 5 submitted prediction attempts. The remaining 6

Feature Set	Model	Station (recent data)			Station (no recent data)		
		+1day	+5days	+7days	+1day	+5days	+7days
FS1	MLP	20.51	27.77	26.75	45.46	48.75	46.06
FS1 (run 1)	LSTM	19.94	25.84	<u>25.42</u>	44.15	46.87	46.38
FS2 (run 3)	LSTM	<u>18.94</u>	25.52	24.95	43.98	46.66	45.36
FS2 (run 4)	TemAtt	18.57	<u>24.96</u>	27.73	44.36	46.60	48.67
FS2	LSTM-GNN	23.84	25.99	30.48	39.84	40.21	44.48
FS2 (run 2)	SpaTemAtt	19.93	24.91	29.92	30.93	34.42	34.85
FS3 (run 5)	SpaTemAtt	20.43	25.18	31.75	<u>35.21</u>	<u>38.46</u>	<u>40.32</u>

Table 1
Best Performing Model Configurations

columns, which present RMSE scores, are split according to those stations that continued to provide data, “Station (recent data)”, and those that stopped supplying data due to malfunctions, “Station (no recent data)” at the station. For the latter, it was necessary to make predictions for these stations based on recent data at other stations. For both station categories, we provide our accuracy for 1-day ahead, 5-days ahead and 7-days ahead predictions. Validation uses root mean squared error, so the lower the error cost, the better the model.

Analysis. Stations with recent data far out-performed those without so we will focus our discussion on the former. As expected, the predictive accuracy for all models decreases as the forecast window increased, significantly for 5 and 7 day predictions. Feature set 2 (FS2) was the best performing across all time intervals. This highlighted the benefit of our gap filling method but also indicated that the camera data did not deliver any added performance. In terms of models, no single model performed best although our Temporal Graph model (TemAtt) was the best-performing model for 1-day predictions and the Spatio-Temporal Graph model (SpaTempAtt) was best performing for 5-day predictions. Finally, all test models, including both of our Graph models, showed significant degradation in performance as sensitivity decreases.

4. Conclusions and Lessons Learned

Some high-level conclusions from our results show that our gap filled data have a positive impact on the predictions, it approximates missing values better and adds spatial information from other locations. Time series specialised models are better at predictions in short-term as LSTM and TemAtt architectures out-perform MLP and simpler models marginally.

While the results provide answers to some questions, there are many other unanswered ones. For example it is not clear why the TemAtt model outperformed the SpaTemAtt model. It may be because the inclusion of spatial info made SpaTemAtt weaker somehow. We were surprised by the comparatively poor performance of the LSTM-GNN. For us, the CCTV images added no value as FS3 featured outside the top-2 across all prediction intervals possibly because there is little overlap between CCTV and air quality data thus little for the models to train on. One other reason for this could be that additional traffic features do not really provide new information. Even if we assume air quality is affected by traffic, was our vehicle counting too naive or is there another explanation.

Acknowledgements

This work was supported in part by Science Foundation Ireland through the the Insight Centre for Data Analytics (SFI/12/RC/2289_P2) and the Centre for Research Training in Machine Learning (18/CRT/6183). We thank the organisers for running the task.

References

- [1] M.-S. Dao, T.-H. Dang, T.-L. Nguyen-Tai, T.-B. Nguyen, D.-T. Dang-Nguyen, Overview of MediaEval 2022 Urban Air: Urban Life and Air Pollution, in: Proc. of the MediaEval 2022 Workshop, 13-15 January 2023, Bergen, Norway and Online, 2022.
- [2] R. W. Pinder, J. M. Klopp, G. Kleiman, G. S. Hagler, Y. Awe, S. Terry, Opportunities and challenges for filling the air quality data gap in low- and middle-income countries, *Atmospheric Environment* 215 (2019) 116794. URL: <https://www.sciencedirect.com/science/article/pii/S1352231019304236>. doi:<https://doi.org/10.1016/j.atmosenv.2019.06.032>.
- [3] E. Falge, D. Baldocchi, R. Olson, P. Anthoni, M. Aubinet, C. Bernhofer, G. Burba, R. Ceulemans, R. Clement, H. Dolman, A. Granier, P. Gross, T. Grünwald, D. Hollinger, N.-O. Jensen, G. Katul, P. Keronen, A. Kowalski, C. T. Lai, B. E. Law, T. Meyers, J. Moncrieff, E. Moors, J. W. Munger, K. Pilegaard, Ü. Rannik, C. Rebmann, A. Suyker, J. Tenhunen, K. Tu, S. Verma, T. Vesala, K. Wilson, S. Wofsy, Gap filling strategies for defensible annual sums of net ecosystem exchange, *Agric. For. Meteorol.* 107 (2001) 43–69. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0168192300002252>. doi:10.1016/S0168-1923(00)00225-2.
- [4] D. Hui, S. Wan, B. Su, G. Katul, R. Monson, Y. Luo, Gap-filling missing data in eddy covariance measurements using multiple imputation (MI) for annual estimations, *Agric. For. Meteorol.* 121 (2004) 93–111. doi:10.1016/S0168-1923(03)00158-8.
- [5] A. M. Moffat, D. Papale, M. Reichstein, D. Y. Hollinger, A. D. Richardson, A. G. Barr, C. Beckstein, B. H. Braswell, G. Churkina, A. R. Desai, E. Falge, J. H. Gove, M. Heimann, D. Hui, A. J. Jarvis, J. Kattge, A. Noormets, V. J. Stauch, Comprehensive comparison of gap-filling techniques for eddy covariance net carbon fluxes, *Agric. For. Meteorol.* 147 (2007) 209–232. URL: <http://linkinghub.elsevier.com/retrieve/pii/S016819230700216X>. doi:10.1016/j.agrformet.2007.08.011.
- [6] Y. Kim, M. S. Johnson, S. H. Knox, T. A. Black, H. J. Dalmagro, M. Kang, J. Kim, D. Baldocchi, Gap-filling approaches for eddy covariance methane fluxes: A comparison of three machine learning algorithms and a traditional method with principal component analysis, *Glob. Chang. Biol.* 26 (2020) 1499–1518. URL: <https://doi.org/10.1111/gcb.14845>. doi:<https://doi.org/10.1111/gcb.14845>.
- [7] C. Li, L. Li, H. Jiang, K. Weng, Y. Geng, L. Li, Z. Ke, Q. Li, M. Cheng, W. Nie, Y. Li, B. Zhang, Y. Liang, L. Zhou, X. Xu, X. Chu, X. Wei, X. Wei, YOLOv6: A single-stage object detection framework for industrial applications, 2022. URL: <https://arxiv.org/abs/2209.02976>. doi:10.48550/ARXIV.2209.02976.
- [8] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C. L. Zitnick, Microsoft COCO: Common objects in context, in: *European Conference on Computer Vision (ECCV)*, Springer, 2014, pp. 740–755.
- [9] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, T.-Y. Liu, Lightgbm: A highly efficient gradient boosting decision tree, *Advances in Neural Information Processing Systems* 30 (2017) 3146–3154.
- [10] T. Chen, C. Guestrin, Xgboost: A scalable tree boosting system, in: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 785–794.
- [11] Technical assistance document for the reporting of daily air quality – the air quality index, <https://www.airnow.gov/sites/default/files/2020-05/aqi-technical-assistance-document-sept2018.pdf>, 2018. U.S. Environmental Protection Agency, EPA 454/B-18-007.
- [12] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Comput.* 9 (1997) 1735–1780.
- [13] B. Yu, H. Yin, Z. Zhu, Spatio-temporal graph convolutional neural network: A deep learning framework for traffic forecasting, *CoRR abs/1709.04875* (2017). URL: <http://arxiv.org/abs/1709.04875>. arXiv:1709.04875.