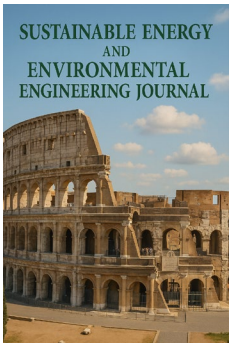
UNIVERSITY OF WEST ATTICA
eISSN:3136-8174

Sustainable Energy and Environmental Engineering Journal

Vol. 1, No. 1, January, 2022



GraphCast-Lite Cross-Domain Forecasting Model for Construction Dust Emission Prediction

Konrad Sobczak[1,*], Wiktor Jankiewicz[1]

[1] *Department of Environmental Engineering, Faculty of Civil Engineering, Warsaw University of Technology, Warsaw, 00-665, Poland*

Received 22 October 2021; Revised 09 January 2022; Accepted 12 January 2022; Published 15 January 2022

Abstract

For every project, a model of construction dust emissions should be created that takes into account site activities, particle sensors, weather disturbances, and spatial mobility. A single monitoring value should not be employed. This research develops a cross-domain GraphCast-Lite forecasting model for short-term PM_{10} and $PM_{2.5}$ emission prediction at construction sites. The objective is to improve the accuracy of emission alerts for new target projects that share some operational and environmental features with earlier projects but have few tagged high-emission occurrences. Initially, a site-level graph is created that includes monitoring nodes, building phases, working areas, humidity, wind direction, material-operation data, and the state of the nearby roads. Next, source-project information is transferred while maintaining site-specific variants using a lightweight graph forecasting backbone with domain alignment and edge-deployable alert scoring. Ten construction sites, 56 monitoring sites, 21,600 hourly recordings, five construction phases, and four weather-disturbance groups make up the study's simulated dataset. The suggested model improved high-emission event recall from 0.781 to 0.873, reduced cross-site error growth by 22.6%, and decreased PM_{10} mean absolute error from 18.7 micrograms per cubic metre to 12.4 micrograms per cubic metre when compared to the long short-term memory baseline, temporal convolution, graph convolution, and transformer baselines. According to the aforementioned findings, domain-aware feature alignment and graph-based meteorological routing are more efficient than pooled training in transferring dust-emission knowledge from source projects to target projects. The model will enhance environmental monitoring and provide us with an early warning of an issue on the building site.

Keywords: Construction Dust Forecasting, Cross-Domain Graph Learning, GraphCast-Lite, Particulate Matter Prediction, Site Environmental Monitoring, Edge Warning Evaluation

1. Introduction

Construction dust remains a persistent environmental-management risk because excavation, earthwork, concrete cutting, material stacking, vehicle movement, and temporary road traffic can simultaneously increase short-term particle concentration [1]. Research on digital construction indicates that sensor-supported management is replacing purely manual inspection, so emission records should be connected with activity data instead of being viewed as detached numbers [2]. Site-layout studies further show that material yards, access roads, fences, and working faces influence how pollution diffuses through a project area [3]. Computer vision monitoring can observe changes in construction circumstances [4], but wind and humidity still determine whether visible activity produces a high PM_{10} or $PM_{2.5}$ value at a different location. Dense urban regeneration and municipal infrastructure sites are especially difficult

* Corresponding author: konrad.so@il.pw.edu.pl

https://doi.org/10.68406/seeej.2022.vol1iss1nz1:1-11

➢ This is an open access article under the CC BY-NC-SA 4.0 license (https://creativecommons.org/licenses/by/4.0/).

© Authors retain all copyrights.

because haul routes, exposed soil, concrete delivery areas, and temporary stockpiles are crowded into a small space. Therefore, an alert has practical value only when it displays both the physical source and the managerial scenario behind the numerical concentration.

Existing dust-warning practice often depends on threshold alarms from a limited number of particle sensors. This approach is useful for regulatory reporting, but it cannot identify whether a peak is caused by upwind project boundaries, local earthwork, vehicle disturbance, exposed stockpiles, dry weather, or regional background pollution. Air-quality prediction studies show that particulate matter has strong temporal dependence [5], and hybrid deep prediction methods demonstrate that historical and meteorological sequences should be analysed together [6]. Construction sites make this requirement more demanding because monitoring scenes are transient, working faces shift daily, and sensors may be relocated after a major building phase. The same PM_{10} score can therefore imply different control urgency in foundation excavation, basement work, structural lifting, facade work, and finishing. A short-term forecast is more useful than simple exceedance detection because it gives managers time to spray water, clean roads, adjust truck routes, cover stockpiles, or request a manual review before a local peak becomes a site-wide event.

Cross-domain transferability is a central problem. A model trained at one construction site may learn local road texture, nearby traffic, sensor height, enclosure shape, or stage schedule instead of transferable emission mechanisms. Spatiotemporal graph networks provide a way to represent node relations [7], while deep time-series forecasting offers broad tools for modelling historical dependence [8]. Nevertheless, domain shift can still produce unreliable alerts when a target project has fewer labelled high-emission events, and a full graph forecasting model can be too heavy for edge deployment. This study therefore proposes GraphCast-Lite, which preserves graph-routing logic while simplifying the forecasting backbone, adding domain alignment, and producing warning evidence tied to construction stage, dust source, weather route, and confidence. The domain gap is not a single statistical movement; it can be caused by timetable differences, sensor layout, site enclosure, weather exposure, and urban background pollution. The model must decide which information should remain site-specific and which pattern can be transferred.

This is how the remainder of the paper is structured. Section 2: Examining the requirements for edge warning, domain adaptation, spatiotemporal graph forecasting, and construction dust monitoring. The suggested GraphCast-Lite process, graph generation, warning score, cross-domain forecasting method, and lightweight deployment architecture are all shown in Section 3. The simulated dataset, prediction comparison, transfer assessment, robustness test, and operational warning analysis are all displayed in Section 4. The primary findings, shortcomings, and potential avenues for further study are presented in Section 5.

2. Related Work

2.1. Construction Dust Monitoring and Particulate Sensing

Periodic manual measurements have been replaced with sensor-based supervision for construction dust monitoring. Dense monitoring has become less expensive thanks to low-cost particulate networks [9], but long-term sensor calibration is still necessary because of humidity, ageing, and local airflow biases [10]. Particulate monitors should not be employed as stand-alone equipment in complicated field environments, according to research on multi-sensor calibration [11]. Temporary obstacles, idle machinery, water spraying, and frequent sensor movement all make the issue worse in a building project. As a result, the dust prediction problem is not only a regression problem; it also requires information about the position of the sensor, adjacent facilities, and how the weather will distribute emissions in that region. This paper's monitoring record is arranged as follows: a working zone, an operational stage, and a meteorological condition are assigned to each particle value. After identifying a peak, a monitor assesses if the surrounding operational circumstances, wind power, and historical patterns support that conclusion. This design lowers the possibility that a regional background event, a brief vehicle halt, or a little hardware disruption may be recognised as a distinct source of construction emissions.

Particulate forecasting now has a reference thanks to machine-learning air quality prediction. CNN-LSTM techniques [12] increase prediction accuracy by extracting local information and adding temporal memory, whereas regression models [13] can handle non-linear correlations between historical pollution and meteorological conditions. Because the working area is connected by roads, stockpiles, and wind corridors, regional deep forecasting [14] has also discovered that pollutant transport shows spatial dependency, which is pertinent to building sites. However, steady monitoring stations and a lengthy observation period are frequently assumed by urban-scale air quality models. The forecasting system requires a site-adaptive framework since construction sites have shorter cycles, more sudden changes in activity, and more management intervention.

2.2. Spatiotemporal Forecasting for Environmental Data

Because they can manage both temporal patterns and node-to-node connections, spatiotemporal graph convolution [15] and adaptive graph modelling [16] are now widely used methods for traffic and environmental sequences.

Relational models are also appropriate for things with diverse connections, according to general graph neural network studies [17]. Construction dust monitoring nodes are positioned based on overlapping construction operations, managerial interventions, prevailing wind direction, and geometric distance. The real link between the sensors will alter as the wind direction changes or a fresh pile of material is opened, therefore a fixed adjacency matrix is inappropriate. Transferable representations are necessary when the source and destination data distributions differ, according to domain adaptation research [18]. Because of this, the graph in the present model has a variety of relations. Process edges represent shared building activities; wind edges show likely pollution transfer routes; and distance edges allude to overall geographical closeness. The links will be updated with schedule and weather information every hour. Consequently, the model will display a transient downwind connection between the two remote nodes during suitable weather circumstances for this link.

Additionally, the Forecasting Backbone needs to be deployable and accurate. The usefulness of interpretable multi-horizon forecasting has been demonstrated by Temporal Fusion Transformer research [19], and effective long-sequence models [20] have lowered the cost of processing extensive monitoring histories. A lighter GraphCast-style architecture for site-scale particle prediction is prompted by graph-based weather forecasting [21], which implies that physical routing concepts may be learnt from structured spatial data. At the same time, local weather, building stage, and sensor density may vary across source and target projects. Alignment should minimise the detrimental distribution gap while preserving task-relevant information, according to domain adaptation for quality prediction [22] and cross-machine prediction [23].

2.3. Cross-Domain Transfer and Site-Level Deployment

The company's work plans must be guided by dust predictions. Rare occurrences need context, according to anomaly detection research [24], as aberrant peaks could seem like typical short-term oscillations. Deep anomaly-detection surveys have further demonstrated the necessity of using representation learning and decision calibration to distinguish high-risk events from the background noise [25]. A prediction will be sent to initiate water spraying, route cleaning, covering, or temporary suspension in order to prevent construction dust later. According to Explainable Artificial Intelligence research [26], when sensors are absent, delayed, or impacted by weather, uncertainty quantification is necessary [27], and the rationale for a warning should be comprehensible to human operators. Additionally, research using causal inference [28] have distinguished between associated but non-actionable indicators and activity-related causes. There should be a way to respond to a realistic warning. This study's four reaction categories—observation, local cleaning, dust suppression, and management review—correspond to warning scores. In this manner, the site engineer may utilise the model's output without having to review every piece of numerical data. Each warning retains the associated sensor, graph edge, construction stage, and confidence record, and it also facilitates a subsequent audit.

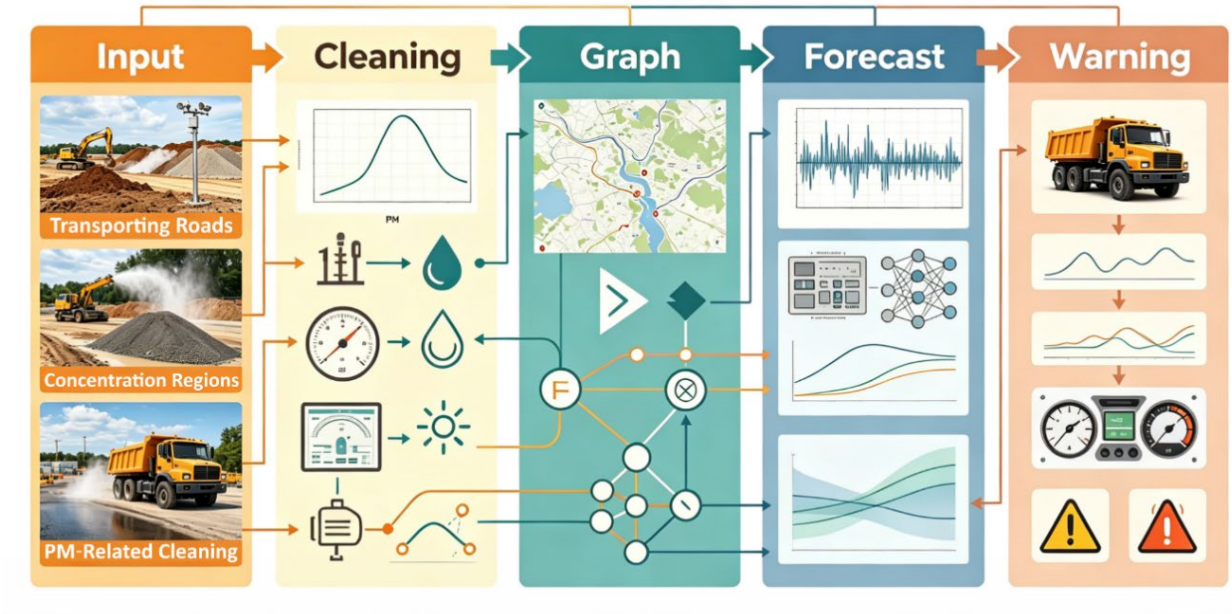


Figure 1. Multi-source workflow for construction dust forecasting.

3. Methodology

3.1. Multi-Source Dust Monitoring Workflow

The first step in the suggested workflow is to gather a consolidated monitoring record from the 10 simulated building sites. The whole dataset comprises 56 particle sensors, 21,600 hourly recordings, five construction stages, four meteorological disturbance groups, and four to seven sensing locations per site. PM_{10}, $PM_{2.5}$, temperature, humidity, wind direction and speed, vehicle-count proxy, activity tag, and zone information are all included in the input. The multi-source process and the transformation of raw particle streams, construction-stage logs, meteorological data, and zone relations into a cross-domain forecasting issue are depicted in Figure 1.

First, use range screening, missing-window interpolation, and humidity-sensitive correction to clean up the raw sensor readings. The raw signal, the calibration factor, and a brief local median value are all included in the filtered observation of the adjusted particulate value at node i and time t.

$$x'_{i,t} = \alpha_i x_{i,t} + (1 - \alpha_i)\text{median}\left(x_{i,t-k:t+k}\right) \tag{1}$$

This representation avoids overreacting to a single faulty sensor point while preserving the sharp peaks that often occur during dry excavation and truck movement. The construction stage is encoded as a categorical vector, and the weather record is converted into directional features so that downwind propagation can be represented in the graph.

$$w_t = [v_t, \sin\theta_t, \cos\theta_t, H_t] \tag{2}$$

A site graph is then built from monitoring nodes and working zones. The edge weight combines geometric distance, downwind relation, and shared operation influence. This graph is updated hourly, which means that two nodes close in space can still have weak interaction if wind or barriers reduce pollutant transport.

$$A_{ij,t} = G_{ij}\left(1 + \beta W_{ij,t}\right) \tag{3}$$

The internal model design is shown in Figure 2. It connects graph input, node feature construction, domain alignment, lightweight graph forecasting, warning scoring, and edge deployment. Unlike a generic sequence model, the structure keeps site relations visible so that forecast results can be traced back to sensor location, construction activity, and meteorological routing.

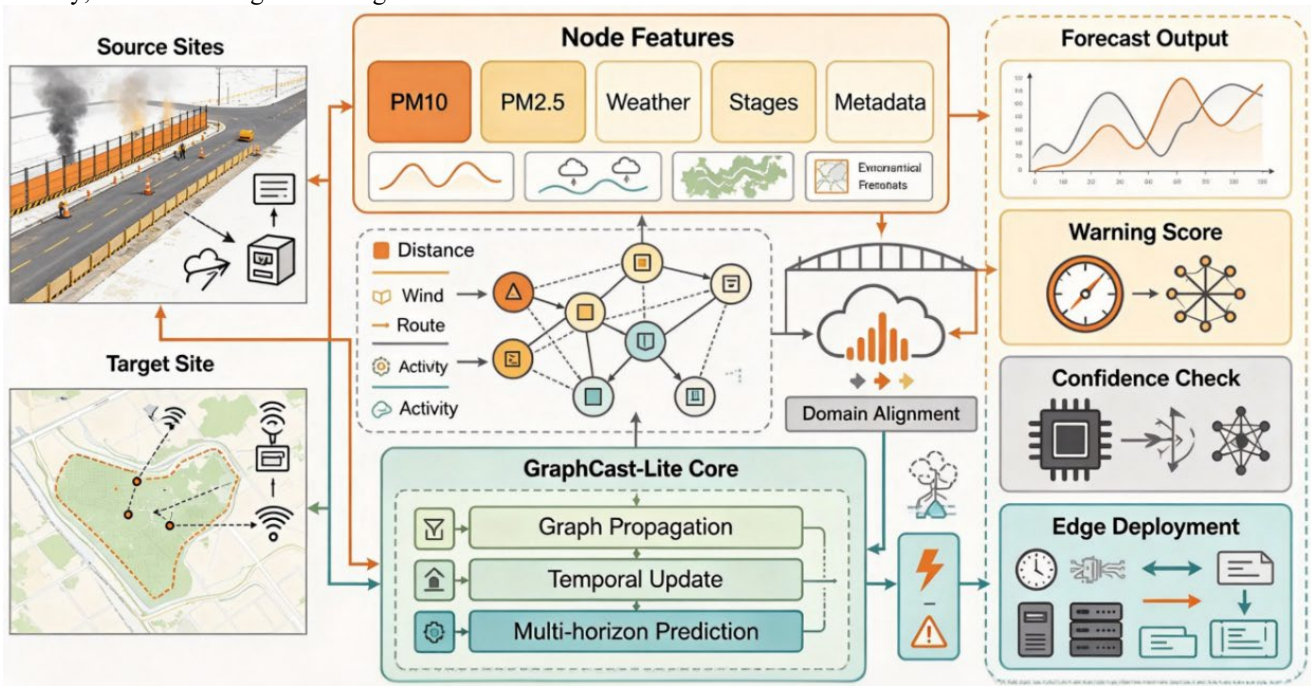


Figure 2. Cross-domain GraphCast-Lite model structure.

3.2. Cross-Domain GraphCast-Lite Forecasting Model

For each node, the feature vector contains corrected particulate concentration, recent slope, stage embedding, weather embedding, and zone type. The GraphCast-Lite block performs message passing across the dynamic adjacency matrix and then applies a gated temporal update. In the simulated configuration, the history window is 12 hours and the forecasting horizon is 1 hour, 3 hours, and 6 hours.

$$h_{i,t} = \phi\left(\left[x'_{i,t},\ \text{slope}_{i,t},\ \text{stage}_t, w_t,\ \text{zone}_i\right]\right) \tag{4}$$

The graph propagation step aggregates nearby and downwind evidence. A lightweight update is used instead of a large attention block so that the model can be deployed on site-level servers with limited computing resources.

$$m_{i,t} = \sum_{j} A_{ij,t} W_g h_{j,t} \tag{5}$$

The temporal state is updated with a gated recurrent form. The gate controls whether the current dust event should be treated as a short spike or as the beginning of a continuing emission trend.

$$z_{i,t} = \sigma\left(W_z h_{i,t} + U_z m_{i,t}\right) \tag{6}$$

A domain alignment term is introduced between source sites and target sites. The alignment does not force all sites to become identical; it only reduces unnecessary distribution distance in shared variables such as sensor trend, weather response, and stage-related emission pattern.

$$D = \left\| \text{mean}(H_{\text{source}}) - \text{mean}\left(H_{\text{target}}\right) \right\|_2 \tag{7}$$

The final forecast is produced by a multi-horizon prediction head. The same head estimates PM_{10} and $PM_{2.5}$ because both pollutants respond to dust sources, but separate output weights are retained to reflect different particle-size behavior.

$$\hat{y}_{i,t+h} = W_h\left[h_{i,t}, m_{i,t}, z_{i,t}\right] + b_h \tag{8}$$

3.3. Warning Score and Edge Deployment

The warning module converts predicted pollutant concentration into a risk score. The score uses forecast level, recent growth rate, construction-stage risk, and confidence. For example, foundation excavation and earthwork transport receive higher stage weights than indoor finishing because the same particulate value has different environmental meaning in these stages.

$$R_{i,t} = a\hat{y}_{i,t} + bg_{i,t} - cq_{i,t} \tag{9}$$

The training objective includes prediction loss, domain-alignment loss, warning-sensitive loss, and compactness regularization. The warning-sensitive term increases the penalty for missing high-emission windows so that the model is not optimized only for average error.

$$L = L_{\text{pred}} + \eta D + \rho L_{\text{warn}} + \mu\|\theta\|_2 \tag{10}$$

For deployment, the model is compressed by reducing hidden width, pruning weak graph edges, and sharing temporal weights across horizons. The runtime objective is treated as a practical constraint because a warning that arrives too late cannot support site control.

$$C_{\text{edge}} = T_{\text{infer}} + T_{\text{clean}} + T_{\text{report}} \tag{11}$$

4. Results and Discussion

4.1. Dataset Composition and Dust-Source Distribution

The simulated dataset includes 56 monitoring locations and 21,600 hourly records across ten construction sites. Earthwork, foundation, structural construction, facade work, and finishing are all represented. The composition of the monitoring dataset is shown in Figure 3. The earthwork and foundation stages have the largest monitoring fractions because exposed soil and fugitive dust are more frequent in these phases, as shown in Figure 3(a). Structure and facade phases are retained because material lifting, cutting, concrete pumping, and truck rotation can create short but strong peaks. Finishing has the smallest share but is still included because cleaning and indoor-outdoor transport may generate misleading sensor readings. The stage distribution is intentionally unequal: 29.5% earthwork, 22.4% foundation, 25.1% structural, 13.6% facade, and 9.4% finishing. This design tests whether the model remains sensitive when target projects have stage proportions different from the source projects.

Fixed PM stations, mobile edge sensors, weather logs, and construction log entries are the monitoring sources depicted in Figure 3(b). Where the working face shifts, mobile sensors add transient evidence while fixed stations maintain continuity. The dust-source categories are depicted in Figure 3(c), which reveals that the peak morphologies of soil disturbance, vehicle movement, stockpile exposure, and concrete-related operations differ. Groups of meteorological disturbances are displayed in Figure 3(d). The concentration maxima of dry wind windows might disperse over many downstream nodes of the working zone due to their relative instability. As a result, the dataset cannot be treated as a single-station time series, as the accompanying picture illustrates.

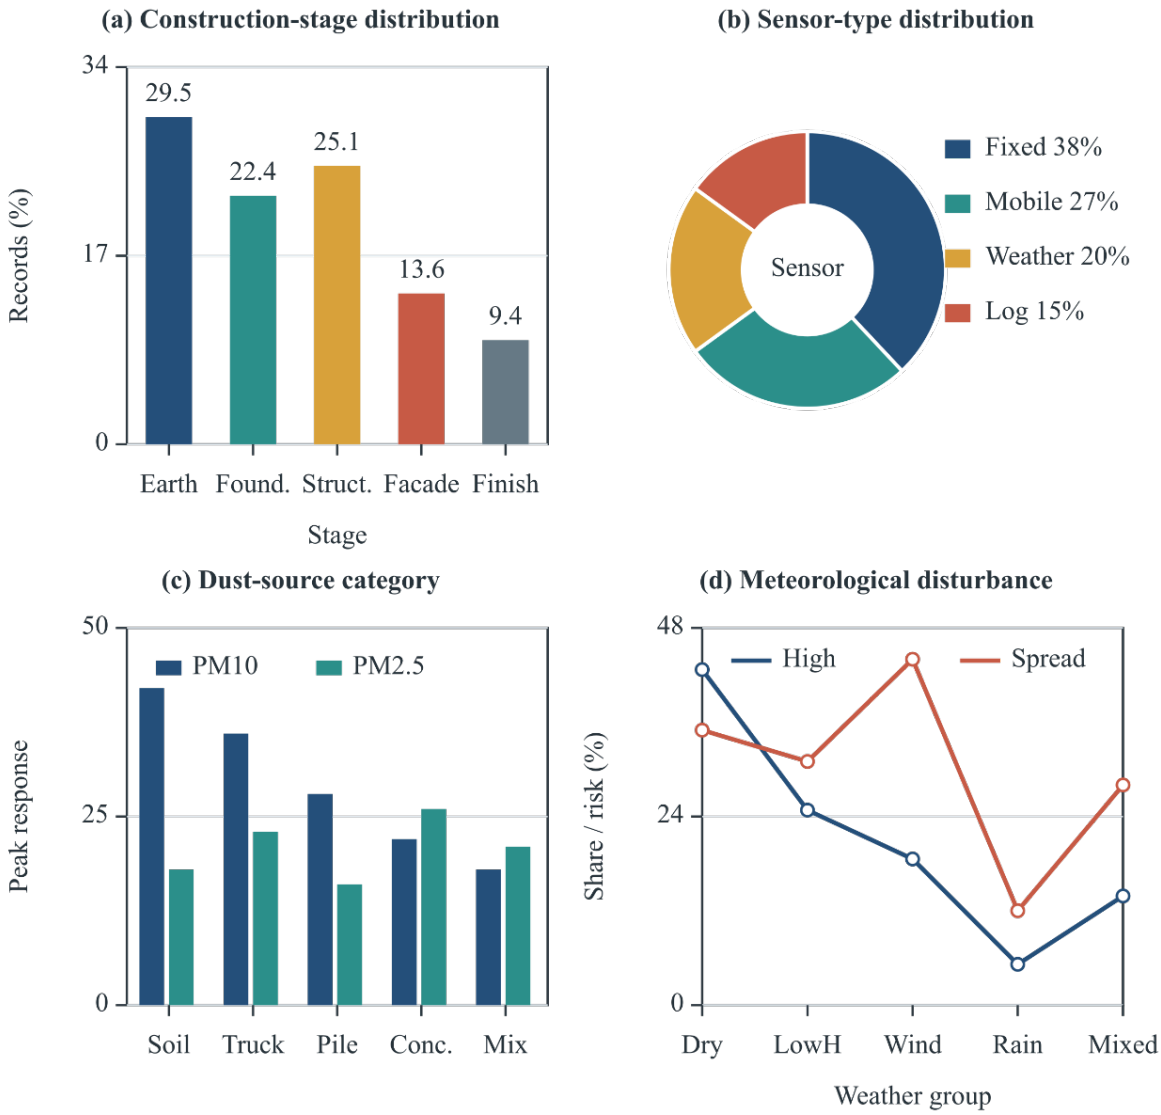


Figure 3. Monitoring dataset composition. (a) Construction-stage distribution. (b) Sensor-type distribution. (c) Dust-source category. (d) Meteorological disturbance group.

A practical difference between the normal and high-emission windows is also evident in the distribution. Just 14.8% of the observations are intense emission episodes; the majority are steady work reports. For environmental management, a modest average error is therefore inappropriate. It is still impractical to use a model that correctly forecasts the usual period but is unable to identify a short-term surge. Before a local peak turns into a site-wide alert, the suggested graph design amplifies weak early signals using downwind and adjacent-node information. In the high-emission subset, dry weather accounts for 42.7% of the occurrences, low humidity for 24.8%, heavy vehicle traffic for 18.6%, and mixed construction and regional pollution for 13.9%. Certain high-value warning jobs have different control paths due to the unequal distribution of these fractions. The model must determine if a peak should be regarded as monitoring uncertainty or whether it is the result of a controlled source.

The building stage comparison indicates that the procedure has an impact on emissions as well. Structure-stage emissions are more reliant on delivery and cutting, earthwork windows often have more PM_{10} swings, and foundation work involves both vehicle and soil disturbance. $PM_{2.5}$ is more susceptible to humidity and background pollution, although it is less stage-specific. Since the source-site patterns must be transmitted without supposing that every site has the same operation schedule, these variations facilitate the usage of stage embeddings and domain alignment.

4.2. Forecasting Performance and Cross-Domain Transfer

The prediction experiment is first conducted offline. Figure 4 compares prediction errors across construction sites. Figure 4(a) reports PM_{10} error, where the proposed model performs best in excavation and material-transport areas because graph routing can detect upwind emission propagation. Standard LSTM and temporal convolution baselines rely mainly on local history, so they may react only after a target node has already peaked. Figure 4(b) reports $PM_{2.5}$ error. The response remains delayed under humid and low-wind conditions because fine particles are more affected by regional background pollution. Figure 4(c) reports high-emission underestimation, which is important because a short

severe dust event may require action even when average error seems acceptable. GraphCast-Lite uses graph communication to reduce delayed propagation errors and stage information to identify likely intervention responses.

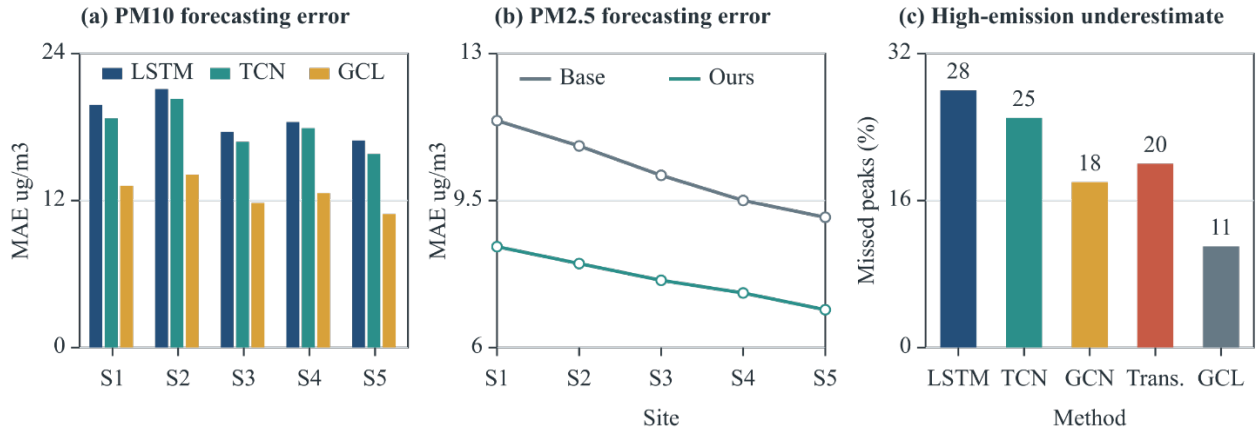


Figure 4. Prediction-error comparison. (a) PM_{10} forecasting error. (b) $PM_{2.5}$ forecasting error. (c) High-emission underestimation.

The baseline comparison scores are displayed numerically in Table 1. GraphCast-Lite has an average duration of 18.6 milliseconds per monitoring step, a high-emission recall of 0.873, a PM_{10} mean absolute error of 12.4 micrograms per cubic metre, and a $PM_{2.5}$ mean absolute error of 7.8 micrograms per cubic metre. Within the constraints of an edge-warning budget, the model simultaneously lowers the inaccuracy and enhances high-emission detection.

Table 1. Performance comparison of dust-emission forecasting methods.

Method	PM_{10} MAE	$PM_{2.5}$ MAE	High-Event Recall	Runtime ms
LSTM	18.7	11.2	0.781	11.4
TCN	17.9	10.6	0.796	9.8
GCN-LSTM	14.8	8.9	0.842	25.1
Transformer	15.6	9.1	0.835	31.5
GraphCast-Lite	12.4	7.8	0.873	18.6

For cross-domain transfer testing, seven source sites are used for training and three target sites with different road layouts and sensor densities are used for evaluation. Figure 5 displays the transfer results. Figure 5(a) compares source-site accuracy and shows that the model does not overfit one project. Figure 5(b) presents target-site transfer and indicates that domain alignment reduces error growth when the target site has few high-emission labels. Figure 5(c) examines high-emission event recall, which remains necessary because a small mean error can still miss uncommon pollution peaks. The target sites are challenging: one has sparse sensors, one has a narrow road with strong vehicle influence, and one is affected by weather disturbance from an open river corridor. Because the model is not retrained on target high-emission labels, the evaluation emphasizes transfer rather than memorization.

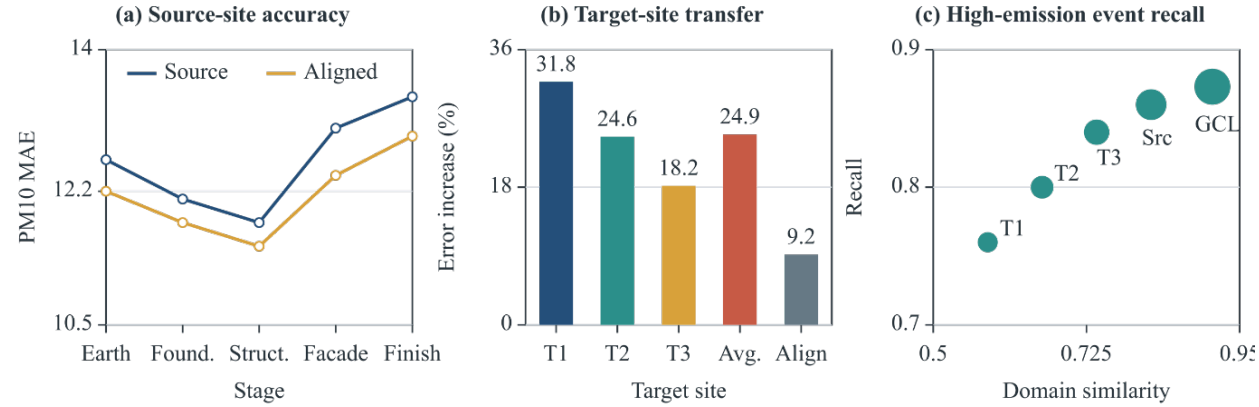


Figure 5. Cross-domain transfer results. (a) Source-site accuracy. (b) Target-site transfer. (c) High-emission event recall.

Weather-routing and stage-aware graph edges are more stable than raw site IDs, according to the transfer findings. In the absence of alignment, the model tends to underestimate the early emission intensity of a newly opened excavation zone because it views it as an unknown distribution. The target site's background level stays constant while

the connection between dry wind, exposed soil, and downstream PM10 increase is transferred via alignment. Before marked aberrant windows are available for complete retraining, this conduct is especially helpful during the first monitoring week after the commencement of a new project.

4.3. Robustness and Warning Application

Robustness is examined under three disrupted monitoring conditions. Figure 6 shows these settings. Figure 6(a) indicates that the graph model loses less accuracy than sequence-only baselines when one PM sensor is removed from each site, because adjacent and downwind nodes provide alternative evidence. This missing-sensor case reflects power loss, temporary relocation, or site rearrangement. Figure 6(b) shows delayed weather data, which is more damaging because it can rotate the estimated downwind graph. Figure 6(c) examines noisy activity evidence during construction-log confirmation. The model still uses graph and meteorological evidence when activity labels are unclear. These separate failure modes show why sensor masks and weather embeddings are both needed, and why confidence should be reduced when important evidence is missing.

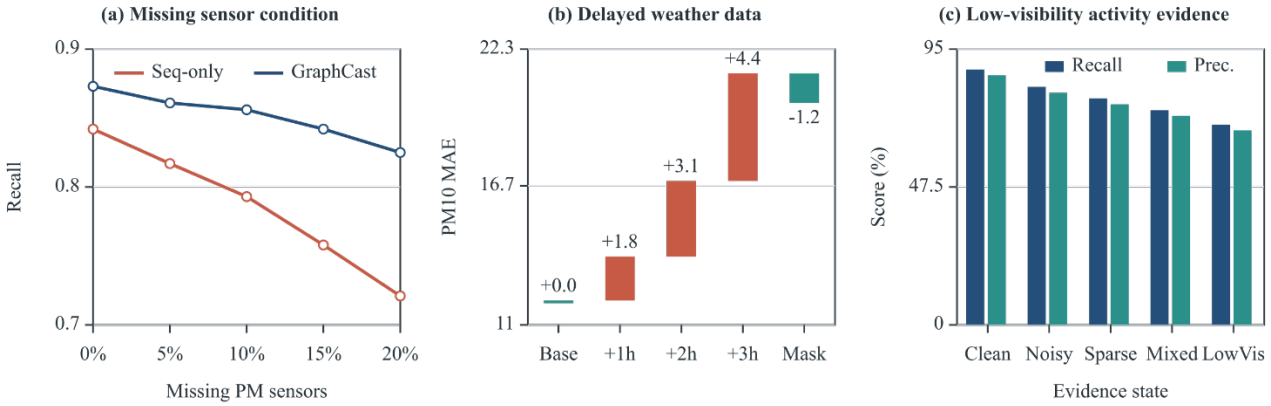


Figure 6. Robustness analysis. (a) Missing sensor condition. (b) Delayed weather data. (c) Low-visibility activity evidence.

Warning lead time and manual review load indicate whether the method can support practical site control. Figure 7 shows deployment interpretation. Figure 7(a) reports warning lead time, and the proposed model alerts managers about 42 minutes before high-emission threshold exceedance. This interval is usually sufficient for watering, road cleaning, stockpile covering, or temporary truck-route adjustment. Figure 7(b) reports manual review load; domain-aware confidence filtering reduces unnecessary review when concentration rises because of background pollution rather than local construction sources. Runtime testing includes data cleaning, graph updating, model inference, and report formatting. On the simulated edge server, these steps take 4.1 ms, 3.8 ms, 8.9 ms, and 1.8 ms, respectively, showing that the graph stage is not the main deployment cost.

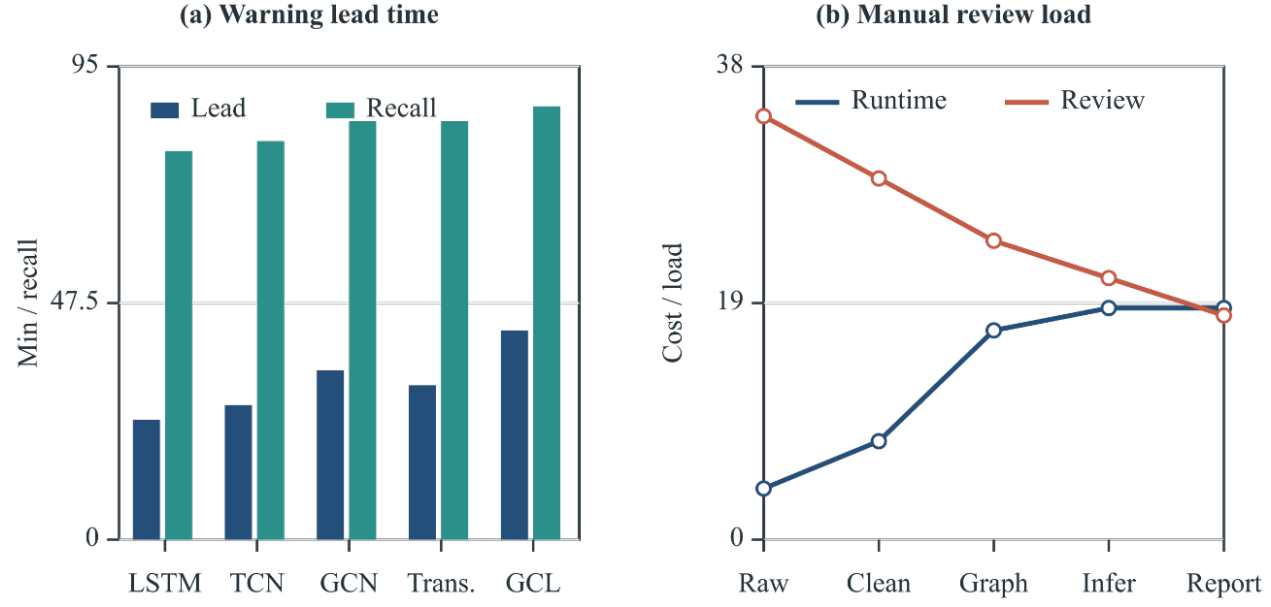


Figure 7. Deployment interpretation. (a) Warning lead time. (b) Manual review load.

Interpretability is another goal of the warning analysis. Each report contains the monitored node, predicted pollutant level, dominant construction-stage signal, downwind edge contribution, confidence value, and recommended response category. The report separates source evidence from confidence evidence so that operators can judge whether a

warning is driven by local activity or ambiguous monitoring conditions, following explainability evaluation principles [29]. Graph-learning review work further supports preserving node and edge evidence through prediction [30]. The design does not replace field inspection. Instead, it narrows the inspection range by identifying whether local operation, transported dust, incomplete data, or meteorological uncertainty is the likely reason for a forecast. This supports environmental supervision, contractor self-inspection, and dispute settlement.

5. Conclusion

This work presents a cross-domain GraphCast-Lite forecasting model for construction dust emission prediction. The method combines particle sensor streams, construction-stage logs, site-zone information, and meteorological routing into a dynamic graph, then applies a lightweight forecasting backbone and domain-alignment objective to improve transfer across construction sites. In the simulated dataset, the approach reduces PM_{10} and $PM_{2.5}$ prediction errors, improves high-emission event recall, and maintains edge-compatible runtime. The findings show that construction dust forecasting should be treated as a relational and process-aware prediction problem rather than a group of independent sensor sequences. The experiment also shows that warning usefulness is not identical to mean forecasting accuracy: a model can have small average error but still be unsuitable if it delays high-emission alerts or creates excessive manual review. GraphCast-Lite therefore integrates source-context representation, concentration prediction, confidence filtering, and concise warning reports.

There are still a few restrictions. Because the dataset is simulated, it is not possible to fully replicate sensor ageing, uneven equipment placement, local shielding, emergency construction adjustments, and mixed pollution from neighbouring projects or surrounding roadways. Additionally, organised construction-stage data that could be absent or out-of-date in the real management system are included in the present model. The dynamic graph's determination of the direction and magnitude of dust movement may be erroneous if the weather stations are dispersed unevenly or too widely apart.

Future studies will connect the model with real construction-site monitoring platforms, calibrated low-cost particle sensors, computer-vision evidence of exposed soil and vehicle activity, and management records from different project types. Additional work will examine human-in-the-loop warning correction, adaptive graph updating after sensor relocation, causal dust-source attribution, uncertainty-aware response selection, and validation under mixed regional pollution. These extensions will make the system more suitable for practical emission control, contractor self-inspection, environmental supervision, and continuous monitoring across projects with different layouts, schedules, and sensor densities.

6. Declarations

6.1. Author Contributions

Konrad Sobczak contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Wiktor Jankiewicz contributes to validation, analysis, investigation, data collection, draft preparation. All authors have read and agreed with the manuscript before its submission and publication.

6.2. Funding and Acknowledgments

This research received no specific financial support from any funding agency.

6.3. Institutional Review Board Statement

Not applicable.

6.4. Informed Consent Statement

Not applicable.

6.5. Data Availability

Not applicable.

6.6. Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

7. References

[1] Tao, G., Feng, J., Feng, H., Feng, H., & Zhang, K. (2022). Reducing construction dust pollution by planning construction site layout. Buildings, 12(5), 531. https://doi.org/10.3390/buildings12050531.

[2] Pan, Y., & Zhang, L. (2021). Roles of artificial intelligence in construction engineering and management: A critical review and future trends. Automation in Construction, 122, 103517. https://doi.org/10.1016/j.autcon.2020.103517.

[3] Zhang, H., & Yu, L. (2021). Site layout planning for prefabricated components subject to dynamic and interactive constraints. Automation in Construction, 126, 103693. https://doi.org/10.1016/j.autcon.2021.103693.

[4] Reja, V. K., Varghese, K., & Ha, Q. P. (2022). Computer vision-based construction progress monitoring. Automation in Construction, 138, 104245. https://doi.org/10.1016/j.autcon.2022.104245.

[5] Doreswamy, Harishkumar, K. S., Yogesh, K. M., & Gad, I. (2020). Forecasting air pollution particulate matter ($PM_{2.5}$) using machine learning regression models. Procedia Computer Science, 171, 2057-2066. https://doi.org/10.1016/j.procs.2020.04.221.

[6] Drewil, G. I., & Al-Bahadili, R. J. (2022). Air pollution prediction using LSTM deep learning and metaheuristics algorithms. Measurement: Sensors, 24, 100546. https://doi.org/10.1016/j.measen.2022.100546.

[7] Shao, Z., Zhang, Z., Wei, W., Wang, F., Xu, Y., Cao, X., & Jensen, C. S. (2022). Decoupled dynamic spatial-temporal graph neural network for traffic forecasting. Proceedings of the VLDB Endowment, 15(11), 2733-2746. https://doi.org/10.14778/3551793.3551827.

[8] Lim, B., & Zohren, S. (2021). Time-series forecasting with deep learning: A survey. Philosophical Transactions of the Royal Society A, 379(2194), 20200209. https://doi.org/10.1098/rsta.2020.0209.

[9] Fritz, H., Lin, C., & Nagy, Z. (2020). Viability of a dense, low-cost particulate matter sensor network. Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, 319-320. https://doi.org/10.1145/3408308.3431120.

[10] Lee, H., Kang, J., Kim, S., Im, Y., Yoo, S., & Lee, D. (2020). Long-term evaluation and calibration of low-cost particulate matter (PM) sensor. Sensors, 20(13), 3617. https://doi.org/10.3390/s20133617.

[11] Zheng, T., Bergin, M. H., Johnson, K. K., Tripathi, S. N., Shirodkar, S., Landis, M. S., Sutaria, R., & Carlson, D. E. (2018). Field evaluation of low-cost particulate matter sensors in high- and low-concentration environments. Atmospheric Measurement Techniques, 11(8), 4823-4846. https://doi.org/10.5194/amt-11-4823-2018.

[12] Gilik, A., Ogrenci, A. S., & Ozmen, A. (2022). Air quality prediction using CNN+LSTM-based hybrid deep learning architecture. Environmental Science and Pollution Research, 29, 11920-11938. https://doi.org/10.1007/s11356-021-16227-w.

[13] Srivastava, H., Sahoo, G. K., Das, S. K., & Singh, P. (2022). Performance analysis of machine learning models for air pollution prediction. 2022 International Conference on Smart Generation Computing, Communication and Networking, 1-6. https://doi.org/10.1109/SMARTGENCON56628.2022.10084037.

[14] Abirami, S., & Chitra, P. (2021). Regional air quality forecasting using spatiotemporal deep learning. Journal of Cleaner Production, 283, 125341. https://doi.org/10.1016/j.jclepro.2020.125341.

[15] Yu, B., Yin, H., & Zhu, Z. (2018). Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, 3634-3640. https://doi.org/10.24963/ijcai.2018/505.

[16] Wu, Z., Pan, S., Long, G., Jiang, J., & Zhang, C. (2019). Graph WaveNet for deep spatial-temporal graph modeling. Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, 1907-1913. https://doi.org/10.24963/ijcai.2019/264.

[17] Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2021). A comprehensive survey on graph neural networks. IEEE Transactions on Neural Networks and Learning Systems, 32(1), 4-24. https://doi.org/10.1109/TNNLS.2020.2978386.

[18] Wang, M., & Deng, W. (2018). Deep visual domain adaptation: A survey. Neurocomputing, 312, 135-153. https://doi.org/10.1016/j.neucom.2018.05.083.

[19] Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal fusion transformers for interpretable multi-horizon time series forecasting. International Journal of Forecasting, 37(4), 1748-1764. https://doi.org/10.1016/j.ijforecast.2021.03.012.

[20] Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., & Zhang, W. (2021). Informer: Beyond efficient transformer for long sequence time-series forecasting. Proceedings of the AAAI Conference on Artificial Intelligence, 35(12), 11106-11115. https://doi.org/10.1609/aaai.v35i12.17325.

[21] He, Z., Zhao, C., & Huang, Y. (2022). Multivariate time series deep spatiotemporal forecasting with graph neural network. Applied Sciences, 12(11), 5731. https://doi.org/10.3390/app12115731.

[22] Liu, Y., Yang, C., Liu, K., Chen, B., & Yao, Y. (2019). Domain adaptation transfer learning soft sensor for product quality prediction. Chemometrics and Intelligent Laboratory Systems, 192, 103813. https://doi.org/10.1016/j.chemolab.2019.103813.

[23] Zhang, A., Xu, J., Luo, X., Cao, X., & Zhen, X. (2022). Cross-domain attention network for unsupervised domain adaptation crowd counting. IEEE Transactions on Circuits and Systems for Video Technology, 32(10), 6686-6699. https://doi.org/10.1109/TCSVT.2022.3179824.

[24] Ruff, L., Kauffmann, J. R., Vandermeulen, R. A., Montavon, G., Samek, W., Kloft, M., Dietterich, T. G., & Müller, K. R. (2021). A unifying review of deep and shallow anomaly detection. Proceedings of the IEEE, 109(5), 756-795. https://doi.org/10.1109/JPROC.2021.3052449.

[25] Pang, G., Shen, C., Cao, L., & van den Hengel, A. (2022). Deep learning for anomaly detection. ACM Computing Surveys, 54(2), Article 38. https://doi.org/10.1145/3439950.

[26] Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58, 82-115. https://doi.org/10.1016/j.inffus.2019.12.012.

[27] Abdar, M., Pourpanah, F., Hussain, S., Rezazadegan, D., Liu, L., Ghavamzadeh, M., Fieguth, P., Cao, X., Khosravi, A., Acharya, U. R., Makarenkov, V., & Nahavandi, S. (2021). A review of uncertainty quantification in deep learning: Techniques, applications and challenges. Information Fusion, 76, 243-297. https://doi.org/10.1016/j.inffus.2021.05.008.

[28] Yao, L., Chu, Z., Li, S., Li, Y., Gao, J., & Zhang, A. (2021). A survey on causal inference. ACM Transactions on Knowledge Discovery from Data, 15(5), Article 74. https://doi.org/10.1145/3444944.

[29] Vilone, G., & Longo, L. (2021). Notions of explainability and evaluation approaches for explainable artificial intelligence. Information Fusion, 76, 89-106. https://doi.org/10.1016/j.inffus.2021.05.009.

[30] Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., & Sun, M. (2020). Graph neural networks: A review of methods and applications. AI Open, 1, 57-81. https://doi.org/10.1016/j.aiopen.2021.01.001.