INDERSCIENCE PUBLISHERS
Linking academia, business and industry through research

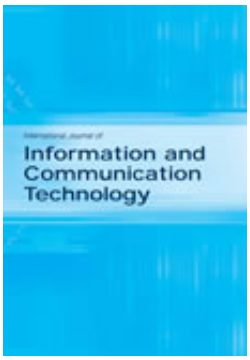

International Journal of Information and Communication Technology

ISSN online: 1741-8070 - ISSN print: 1466-6642
https://www.inderscience.com/ijict

Edge IoT causal graph neural network for enhancing urban carbon sink accounting accuracy

Zhenyu Qian, Qingfeng Zhang, Xiaohong Liu, Youshui Zhang, Zixin Jia

DOI: 10.1504/IJICT.2025.10075268

Article History:
Received: 19 September 2025
Last revised: 17 October 2025
Accepted: 19 October 2025
Published online: 12 January 2026

Copyright © 2025 Inderscience Enterprises Ltd.

Edge IoT causal graph neural network for enhancing urban carbon sink accounting accuracy

Zhenyu Qian and Qingfeng Zhang

College of Natural Resources and Environment,
Northwest A&F University,
Yangling, 712100, China
Email: ZhenyuQian@nwafu.edu.cn
Email: zhangqingfeng@netease.com

Xiaohong Liu* and Youshui Zhang

Shaanxi Forestry Survey and Planning Institute,
Xi'an, 712100, China
Email: zhqf@nwsuaf.edu.cn
Email: 15091625258@163.com
*Corresponding author

Zixin Jia

College of Natural Resources and Environment,
Northwest A&F University,
Yangling, 712100, China
Email: jiazixin@nwafu.edu.cn

Abstract: Global climate change has made accurate urban carbon sink accounting crucial for low-carbon policies and ecological restoration, yet traditional methods suffer from inefficiency, low precision, and poor generalisation. To address these issues, this study proposes an edge IoT-causal graph neural network framework. It integrates a multi-layer edge internet of things architecture reducing data transmission latency by over 40% compared to cloud-centric systems. Additionally, a causal graph neural network model is developed; it infers the causal structure of environmental variables via an improved PC algorithm and embeds this structure into graph attention network training to avoid spurious correlations. Experimental validation on real urban green space data shows the framework achieves 94.7% accounting accuracy, outperforming traditional graph neural networks, support vector machines, and remote sensing inversion by over 8.5%. This work provides a practical technical paradigm for high-precision urban carbon sink accounting, supporting evidence-based urban low-carbon management.

Keywords: edge IoT; causal graph neural network; CGNN; urban carbon sink accounting; causal inference; multi-source data fusion.

Reference to this paper should be made as follows: Qian, Z., Zhang, Q., Liu, X., Zhang, Y. and Jia, Z. (2025) 'Edge IoT causal graph neural network for enhancing urban carbon sink accounting accuracy', *Int. J. Information and Communication Technology*, Vol. 26, No. 51, pp.85–107.

Copyright © The Author(s) 2025. Published by Inderscience Publishers Ltd. This is an Open Access Article distributed under the CC BY-NC-ND license. (http://creativecommons.org/licenses/by-nc-nd/4.0/)

Biographical notes: Zhenyu Qian received his Bachelor's degree from Taiyuan University of Technology in 2023. He is currebtly a Master's student in the College of Natural Resources and Environment, Northwest A&F University. His research interests include carbon sink accounting, soil-water conservation testing.

Qingfeng Zhang received his PhD from Northwest A&F University in 2008. He is currently a Doctoral Supervisor in the Direction of Land Resources and Spatial Information Technology at the Northwest A&F University. His research interests include the utilisation of agricultural water and soil resources, and the use of water-soil resources in spatial information technology.

Xiaohong Liu received her Bachelor's degree from Southwest Forestry University in 2012. She is currently the Director and Engineer at the Shaanxi Forestry Survey and Planning Institute (Shaanxi Forest Resources Monitoring Centre). Her research interests include comprehensive monitoring of forests, measurement of forestry carbon sinks.

Youshui Zhang received his Bachelor's degree from Gansu Agricultural University in 2009. At present, he is a Forestry Engineer at the Shaanxi Forestry Survey and Planning Institute (Shaanxi Forest Resources Monitoring Centre). His research interests are mainly forestry planning consultation, integration of nature reserves and overall planning.

Zixin Jia received his Bachelor's degree from Gansu Agricultural University in 2023. He is currently a Master's student in the College of Natural Resources and Environment, Northwest A&F University. His research interests are mainly land resources utilisation and ecosystem services.

1 Introduction

Global climate change has spurred nations worldwide to commit to carbon peaking and carbon neutrality targets, with cities emerging as pivotal hubs for carbon emission reduction and carbon sink regulation. Urban carbon sinks – encompassing vegetation carbon sinks, soil carbon sinks, and aquatic carbon sinks – can offset roughly 15–20% of urban carbon emissions, rendering accurate carbon sink accounting essential for formulating effective low-carbon policies and assessing the outcomes of ecological restoration efforts (Creutzig et al., 2019). However, the complexity inherent to urban ecosystems – characterised by fragmented green spaces, diverse underlying surfaces, and dynamic environmental shifts – poses substantial hurdles to achieving high-precision, real-time carbon sink accounting (Alshayeb, 2025).

Traditional accounting methods face critical limitations: while field survey methods yield accurate results for small plots, they rely on labour-intensive manual sampling and laboratory analysis such as allometric equations applied to estimate vegetation biomass, leading to prolonged cycles often annual or seasonal that fail to capture short-term carbon sink dynamics triggered by abrupt temperature fluctuations or precipitation events (Dong et al., 2023). Remote sensing inversion methods, which leverage satellite or unmanned aerial vehicle (UAV) data alongside vegetation indices like the normalised difference vegetation index (NDVI) to estimate regional carbon sinks, suffer from low spatial resolution typically ≥ 30 *m* (Wang et al., 2023). This low resolution impedes the

distinction of micro-scale urban green spaces, including street trees and community lawns, while atmospheric interference such as cloud cover further degrades their accuracy. Cloud-centric internet of things (IoT) systems, despite collecting environmental data via distributed sensors, transmit massive volumes of raw data to remote cloud servers – causing transmission latencies to exceed 1 second and straining network bandwidth, thereby hindering real-time carbon sink accounting I (Cohen et al., 2021).

Moreover, conventional machine learning models for carbon sink accounting – including support vector machines (SVMs) and backpropagation (BP) neural networks – prioritise mining statistical correlations between input variables like temperature and soil moisture and carbon sink amounts but overlook the inherent causal mechanisms underlying carbon sink formation (Islam, 2025). For instance, temperature directly drives vegetation photosynthesis a core biological process for carbon sequestration, whereas the observed correlation between atmospheric CO_2 concentration and soil carbon content may be spurious as it is mediated by vegetation growth (Korycki et al., 2025). Such over-reliance on correlational patterns results in poor model generalisation, particularly when faced with shifts in data distribution – such as seasonal vegetation dormancy or extreme weather events.

To address these challenges, edge computing technology deploys computational resources in proximity to sensors to enable real-time pre-processing of raw data including noise reduction and feature extraction, reducing the volume of data transmitted to the cloud. Meanwhile, causal inference techniques – focused on identifying causal relationships rather than mere correlations – enhance model interpretability and robustness. Graph neural networks (GNNs) excel at processing structured data such as the spatial relationships between monitoring points yet lack inherent capabilities for causal modelling (Krich et al., 2022). By combining causal inference with GNNs to form a causal graph neural network (CGNN), it becomes possible to simultaneously capture both the structural characteristics of environmental data and the causal mechanisms governing carbon sink dynamics, thereby laying a foundation for high-precision carbon sink accounting (Kumar et al., 2022).

Driven by these research gaps, this study develops an edge IoT-CGNN framework that integrates edge IoT data acquisition with CGNN to tackle the challenges of high latency, low accuracy, and poor generalisation in urban carbon sink accounting. Specifically, the framework aims to construct a real-time, multi-source edge IoT monitoring system, design a CGNN model capable of learning causal relationships between environmental factors and carbon sink capacity, and verify its performance using real urban monitoring data while comparing it with mainstream accounting methods (Leist et al., 2022).

While prior studies have explored Edge IoT for environmental monitoring or CGNNs in domains like traffic and water quality, this work presents a novel integration specifically tailored for the complex, multi-source data environment of urban carbon sink accounting. Our framework distinguishes itself in three key aspects compared to existing Edge IoT or CGNN studies in other environmental domains. It designs a multi-layer edge architecture specifically for fusing heterogeneous carbon sink data, moving beyond single-source data collection common in existing systems; It develops a causal discovery-driven CGNN model that explicitly embeds an improved PC algorithm into the graph attention network (GAT), shifting the learning paradigm from data-driven correlation to causal mechanism; It establishes a fully functional, real-time accounting pipeline from edge sensing to cloud-based model serving, validated extensively across

diverse urban green spaces, demonstrating both low latency and high accuracy – a combination not yet achieved in prior work.

This work makes four key contributions:

1 A multi-layer edge IoT architecture tailored for urban carbon sink data acquisition is proposed, integrating heterogeneous sensors with edge computing nodes. This architecture enables real-time data pre-processing including noise reduction, drift correction, and multi-source fusion at the network edge, reducing data transmission latency by over 40% compared to cloud-centric systems and alleviating network bandwidth pressure.

2 A causal discovery-driven CGNN model is developed, which first infers the causal structure among environmental variables via constraint-based causal discovery algorithms and then embeds this causal structure into graph neural network training. This design ensures the model captures intrinsic causal mechanisms of carbon sink formation, avoiding spurious correlations and enhancing generalisation across dynamic urban scenarios.

3 The Edge IoT-CGNN framework is comprehensively validated using real-world urban monitoring data from diverse green space types, with results demonstrating that it achieves a carbon sink accounting accuracy of 94.7% – outperforming traditional GNNs, SVMs, and remote sensing inversion methods by 8.5% or more.

4 By bridging edge computing, causal inference, and GNNs, this study provides a practical and generalisable technical paradigm for high-precision urban carbon sink accounting, supporting evidence-based decision-making in urban low-carbon management and ecological restoration strategies.

The remainder of this paper is structured as Section 2 reviews related work on urban carbon sink accounting, edge IoT applications, and causal GNNs to identify research gaps. Section 3 introduces preliminary theories, including edge IoT components, carbon sink accounting indicators, and CGNN basic principles. Section 4 details the edge IoT data acquisition system design, covering architecture, hardware selection, and edge preprocessing algorithms. Section 5 elaborates CGNN model construction, including causal structure discovery and GNN optimisation. Section 6 presents experimental setup, dataset, and results from comparisons with mainstream methods. Section 7 concludes the study and discusses future directions.

2 Relevant work

2.1 SSD urban carbon sink accounting methods

Urban carbon sink accounting has evolved through three core methodological categories, each designed to address specific challenges but facing limitations that restrict their effectiveness in dynamic urban ecosystems. Field survey methods, the most traditional and ground-truthed approach, focus on direct on-site sampling and allometric modelling to calculate carbon storage, with high precision for small plots but poor scalability for large urban areas. For vegetation carbon sinks – the primary component of urban carbon storage – individual plant biomass serves as the foundational input, estimated using

species-specific allometric equations tailored to different plant types. The standard allometric equation for woody plants is:

$$M = a \cdot DBH^b \cdot H^c \tag{1}$$

where M represents the aboveground biomass of a single plant, DBH denotes the diameter of the tree trunk measured 1.3 metres above the ground, H is the total height of the plant, and a, b, c are species-specific calibration coefficients derived from regression analysis of paired 'actual biomass-sample' data – for instance, felling small sample trees of the same species to measure their dry biomass, then fitting the data to the equation to determine coefficients. Once biomass M is obtained, vegetation carbon storage is calculated using:

$$C_{veg} = M \cdot \gamma \tag{2}$$

where C_{veg} is the carbon storage of the vegetation and $\gamma = 0.5$ is the standard carbon content coefficient for terrestrial plants, reflecting the scientific consensus that carbon constitutes approximately 50% of the dry weight of most vascular plants exceptions like succulents have slightly lower coefficients, 0.45–0.48, but 0.5 remains the industry standard for urban vegetation accounting due to its simplicity and wide applicability. Despite their high accuracy for small-scale plots, field surveys suffer from two critical drawbacks: long survey cycles and high labour costs, making them impractical for large urban areas with thousands of hectares of green space (Liu et al., 2024).

To overcome the scalability limitations of field surveys, remote sensing inversion methods leverage satellite or UAV data to estimate regional carbon sinks, using vegetation indices to establish correlations between remote sensing signals and carbon density. The most widely used vegetation index is the NDVI, which quantifies vegetation coverage and vigor by measuring the difference between near-infrared (NIR) light reflectance – vegetation strongly reflects NIR and absorbs red light, while non-vegetated surfaces show the opposite pattern. The formula for NDVI is:

$$NDVI = \frac{\rho_{NIR} - \rho_{Red}}{\rho_{NIR} + \rho_{Red}} \tag{3}$$

where ρ_{NIR} is the reflectance of the NIR band typically 750–900 nm, depending on the satellite sensor and ρ_{Red} is the reflectance of the red band 620–670 nm. NDVI values range from –1 to 1, with values > 0.3 indicating dense vegetation, 0.1–0.3 indicating sparse vegetation, and < 0 indicating non-vegetated surfaces. Based on NDVI, researchers construct linear inversion models to calculate vegetation carbon density:

$$D_{veg} = k \cdot NDVI + d \tag{4}$$

where D_{veg} is the vegetation carbon density, k and d are calibration coefficients derived from ground-truth data, with typical k values ranging from 20–40 t C/(ha·NDVI unit) and d values from –5 to –2 t C/ha for urban green spaces. Commonly used satellite sensors include Landsat 8 with a spatial resolution of 30 metres, suitable for large urban regions and Sentinel-2, but even Sentinel-2's resolution is insufficient for micro-scale urban green spaces – for example, a single street tree with a crown diameter of 5–10 metres may only occupy 1–2 pixels in a Sentinel-2 image, making it impossible to distinguish from adjacent roads or buildings. Additionally, atmospheric interferences such as cloud

cover, aerosols, and haze distort the reflectance values of NIR and Red bands: cloud cover can reduce NDVI by 0.2–0.4, while aerosols scatter light and increase ρ_{Red}, leading to underestimated NDVI values. These interferences reduce inversion accuracy by 15–25% in urban areas, where air pollution and frequent cloud cover are common.

Model-based methods for urban carbon sink accounting include process-based and statistical models, each with distinct strengths and weaknesses. Process-based models simulate the entire carbon cycle of ecosystems by integrating physiological processes and environmental drivers, aiming to capture the mechanistic relationships between environmental factors and carbon sequestration. For example, Biome-BGC calculates gross primary productivity (GPP) as:

$$GPP = PAR \cdot FPAR \cdot \in \tag{5}$$

where PAR is photosynthetically active radiation, $FPAR$ is the proportion of PAR absorbed by vegetation, and $\in$ is the light use efficiency. While process-based models offer high mechanistic accuracy, they require detailed input parameters such as soil texture, vegetation physiological traits, and hourly meteorological data – parameters that are difficult to obtain in urban areas due to the fragmentation of green spaces and the mixing of natural and artificial surfaces. Statistical models, by contrast, simplify carbon sink calculation using empirical coefficients, with the most widely used being the inventory method. The core formula of the method is

$$C = A \cdot D \cdot T \tag{6}$$

where C is the total carbon sink of a specific carbon pool, A is the area of the carbon pool, D is the annual carbon density increment of the pool, and T is the accounting period.

The method is simple to implement and requires minimal data, but it relies on outdated statistical coefficients and cannot capture dynamic changes in urban carbon sinks. In recent years, machine learning models have been applied to carbon sink accounting to improve fitting accuracy, using multi-source data as inputs. However, these models focus on mining statistical correlations between input variables and carbon sink amounts rather than identifying causal relationships, leading to poor generalisation – for example, a random forest model trained to predict carbon sinks in summer may fail in winter, as the correlation between temperature and NDVI changes with the season, and the model cannot distinguish between causal drivers and spurious correlations.

2.2 Edge IoT in environmental monitoring

Edge IoT systems have emerged as a solution to the latency and bandwidth constraints of cloud-centric IoT architectures, with growing applications in environmental monitoring due to their ability to process data in real-time at the network edge. Unlike cloud-centric systems, which transmit all raw sensor data to remote cloud servers for processing leading to high latency and large bandwidth consumption, edge IoT systems deploy edge nodes in proximity to sensors typically within 100–500 metres to perform real-time data pre-processing, reducing the volume of data transmitted to the cloud and lowering latency (Liu et al., 2023). The core workflow of an edge IoT environmental monitoring system involves three steps: first, distributed sensors collect environmental data at fixed intervals; second, edge nodes receive this data via short-range communication protocols and perform pre-processing tasks such as data cleaning, noise reduction, and feature

extraction; third, the pre-processed data rather than raw data is transmitted to the cloud for long-term storage and further analysis, while edge nodes can also generate local alerts for abnormal conditions (Mariappan et al., 2012).

A typical example of edge IoT application in air quality monitoring involves processing PM2.5 data: sensors collect PM2.5 concentrations every 10 seconds, and edge nodes use the 3σ rule to filter abnormal values – an outlier detection method based on the normal distribution, where values outside the range $[\mu–3\sigma, \mu + 3\sigma]$ are classified as abnormal. In this rule, μ is the mean of the PM2.5 concentration over a 5-minute window, and σ is the standard deviation of the 5-minute window (Mohsen et al., 2023). After filtering outliers, edge nodes compute the hourly average PM2.5 concentration and transmit only the hourly averages to the cloud, cutting data transmission by 90% and reducing latency from > 1 second to < 0.3 seconds. Another application is forest ecological monitoring, where edge nodes collect soil moisture data from sensors buried in the ground; when soil moisture falls below a threshold, edge nodes trigger local early warnings for drought stress, eliminating the need to wait for cloud processing and enabling timely forest management.

Despite these successes, edge IoT applications in urban carbon sink monitoring remain limited, with three key shortcomings. First, existing edge IoT systems for environmental monitoring typically collect single-source data rather than integrating the multi-source data required for comprehensive carbon sink accounting – for example, some systems only collect vegetation NDVI data via remote sensing sensors, while others focus solely on soil moisture, but urban carbon sink accounting requires data from four key domains: vegetation, soil, atmosphere, and meteorology. Without integrating these multi-source data, edge nodes cannot provide the comprehensive inputs needed for accurate carbon sink calculation. Second, edge pre-processing algorithms are rudimentary and lack tailored processing for carbon sink-specific data. Most edge nodes only perform basic tasks such as data cleaning and temporal aggregation, but carbon sink data requires specialised pre-processing: for example, CO_2 sensors are prone to drift caused by temperature changes, requiring edge nodes to apply temperature correction using a formula like:

$$C_{corr} = C_{raw} - k \cdot (T - T_0) \tag{7}$$

where C_{corr} is the corrected CO_2 concentration, C_{raw} is the raw concentration from the sensor, T is the current temperature, $T_0 = 25°C$ is the sensor calibration temperature, and $k = 3$ ppm/°C is the temperature drift coefficient – but few existing edge systems include such correction. Additionally, sensors for carbon sink data have disparate sampling frequencies, requiring edge nodes to perform multi-rate data fusion to align time series, but current edge algorithms lack this capability, leading to inconsistent data inputs for carbon sink models. Third, edge nodes rarely participate in carbon sink calculation pipelines, limiting their ability to support real-time accounting. Most edge IoT systems only pre-process data and transmit it to the cloud, where carbon sink calculations are performed – this means real-time carbon sink results still depend on cloud latency even if reduced by pre-processing, and edge nodes cannot generate on-site carbon sink estimates for immediate decision-making.

2.3 Causal graph neural networks

GNNs have become a powerful tool for environmental monitoring tasks due to their ability to model structured data – data with inherent relationships between entities, such as the spatial relationships between urban green space monitoring points (Mohsen, et al., 2023). Traditional GNN variants include graph convolutional networks (GCNs) and GATs, both of which use graph convolutions to aggregate feature information from neighbouring nodes, enabling the model to capture spatial dependencies. The core formula of a GCN layer is

$$\mathbf{H}^{(l+1)} = \widetilde{\mathbf{A}}\mathbf{H}^{(l)}\mathbf{W}^{(l)} \tag{8}$$

where $\mathbf{H}^{(l+1)}$ is the feature matrix of nodes at the $(l + 1)^{\text{th}}$ layer, $\widetilde{\mathbf{A}}$ is the normalised adjacency matrix of the graph, a square matrix where $\widetilde{\mathbf{A}}_{ij} = 1$ if there is a spatial relationship between node i and node j, 0 otherwise, normalised by node degrees to avoid feature scaling issues, $\mathbf{H}^{(l)}$ is the feature matrix at the l^{th} layer, and $\mathbf{W}^{(l)}$ is the weight matrix for the l^{th} layer learned during model training to extract meaningful features. GATs improve on GCNs by introducing an attention mechanism to weight the importance of neighbouring nodes, with the attention coefficient formula:

$$\alpha_{ij} = \frac{\exp\left(\text{LeakyReLU}\left(\mathbf{a}^T\left[\mathbf{W}\mathbf{h}_i \,\|\, \mathbf{W}\mathbf{h}_j\right]\right)\right)}{\sum_{k\in\mathcal{N}(i)} \exp\left(\text{LeakyReLU}\left(\mathbf{a}^T\left[\mathbf{W}\mathbf{h}_i \,\|\, \mathbf{W}\mathbf{h}_k\right]\right)\right)} \tag{9}$$

where α_{ij} is the attention coefficient between node i and node j, $\mathbf{a}$ is the attention vector, $\mathbf{h}_i$ and $\mathbf{h}_j$ are the feature vectors of nodes i and j, $\mathbf{W}$ is the linear transformation matrix, $\|$ denotes vector concatenation, and $N(i)$ is the set of neighbouring nodes of i. While GCNs and GATs excel at capturing spatial correlations, they have a critical limitation: they only model statistical correlations between nodes, not causal relationships – this means they cannot distinguish between direct causal drivers of carbon sinks and spurious correlations. This limitation leads to poor model generalisation: a GAT trained to predict carbon sinks in a temperate city may fail in a tropical city, as the correlational patterns between environmental variables and carbon sinks differ between climates, even if the causal mechanisms remain the same.

CGNNs address this limitation by integrating causal inference with GNNs, combining the structured data modelling capabilities of GNNs with the causal relationship identification of causal inference. The CGNN workflow consists of two core steps: first, learn the causal structure of environmental variables using causal discovery algorithms; second, use this causal structure to guide GNN training, ensuring the model focuses on causal relationships rather than spurious correlations. Causal discovery algorithms aim to construct a directed acyclic graph (DAG) where nodes represent variables and directed edges represent causal relationships. A widely used causal discovery algorithm is the PC algorithm, which tests for conditional independence between variables to remove non-causal edges. The PC algorithm uses mutual information to measure conditional independence: the mutual information between variables X and Y given a set of variables Z is:

$$I(X;Y \mid Z) = \int_{X,Y,Z} p(x,y,z) \log \frac{p(x,y \mid z)}{p(x \mid z)p(y \mid z)} dxdydz \tag{10}$$

where $p(x, y, z)$ is the joint probability density function of X, Y, and Z, $p(x, y|z)$ the conditional joint probability density of X and Y given Z, and $p(x|z)$, $p(y|z)$ are the conditional probability densities of X and Y given Z. If $I(X;Y|Z) < \in$ is a small threshold), X and Y are considered conditionally independent given Z, and the edge between them is removed from the graph. Once the DAG is constructed, the CGNN uses it to modify the GNN aggregation process – for example, in a causal GAT, the attention coefficient a_{ij} is weighted by the causal strength between nodes i and j, resulting in the modified attention coefficient:

$$\alpha_{ij}^{causal} = \alpha_{ij} \cdot S_{ij} \tag{11}$$

where S_{ij} is the causal strength for a direct causal edge and $S_{ij} = 0$ for no causal relationship. This modification ensures the GNN prioritises feature information from variables with direct causal effects on carbon sinks, reducing the influence of spurious correlations.

CGNNs have shown promising results in other environmental and engineering domains. In traffic prediction, CGNNs model causal relationships between traffic flow, weather, and road conditions, improving prediction accuracy by 12–18% compared to traditional GNNs. In water quality prediction, CGNNs capture causal links between hydrological factors and water quality parameters, enabling more accurate predictions of water pollution events even under unseen hydrological conditions. However, to date, no studies have applied CGNNs to urban carbon sink accounting, with three key challenges hindering their adoption. First, constructing a causal DAG that accurately reflects the mechanistic chain of urban carbon sink formation is complex – urban carbon sinks are influenced by a mix of natural and artificial factors, and the causal relationships between these factors are non-linear and context-dependent. This requires a DAG that can capture non-linear causal effects, which traditional causal discovery algorithms like the PC algorithm, which assumes linear relationships struggle to model. Second, integrating heterogeneous multi-source edge IoT data into causal discovery is challenging – edge IoT data for carbon sink accounting includes continuous variables, categorical variables, and count variables, and most causal discovery algorithms are designed for single-type data, requiring specialised pre-processing that can introduce biases if not done carefully. Third, optimising the GNN architecture based on the causal DAG to balance model complexity and computational efficiency is non-trivial – while a more complex GNN can capture detailed causal relationships, it requires more computing resources, which is a problem for edge nodes with limited processing power. Finding a lightweight CGNN architecture that can run on edge nodes while maintaining high accuracy remains an unresolved issue (Necula, 2023).

3 Preliminaries

3.1 Edge IoT technology and urban carbon sink accounting basics

Edge IoT systems designed for urban carbon sink monitoring rely on three interconnected core components to achieve real-time, high-quality data acquisition and preliminary processing. The first component is multi-source sensors, which are deployed across urban green spaces and water bodies to collect data on variables that directly or indirectly influence carbon sink capacity. These sensors capture vegetation-related data including normalised vegetation index, chlorophyll content, soil-related data including soil moisture, soil organic carbon content, atmospheric data including CO_2 concentration, air temperature, air humidity, and meteorological data including precipitation, solar radiation – all of which are essential for comprehensively assessing carbon sink dynamics (Nguyen et al., 2023). The second component is edge nodes, which are equipped with microprocessors such as Raspberry Pi 4 and communication modules including LoRa and NB-IoT. These edge nodes are placed in close proximity to sensor clusters usually within 100–500 metres to perform real-time pre-processing tasks including data cleaning, feature extraction, and data fusion, which reduces the volume of data that needs to be transmitted to the cloud and avoids redundant information. The third component is the cloud platform, which is responsible for storing pre-processed data, training the subsequent CGNN model, and presenting final carbon sink accounting results to users through visual interfaces (Pimenow et al., 2025).

A critical performance metric for edge IoT systems in carbon sink monitoring is data transmission latency, which is defined as the total time from when a sensor collects raw data to when the pre-processed data is successfully stored in the cloud platform. For real-time urban carbon sink accounting – where timely adjustments to green space management or low-carbon policies may be required – this latency must be less than 0.5 seconds. The latency of the edge IoT system is calculated using the formula:

$$T = T_s + T_p + T_t \tag{12}$$

where T represents the total data transmission latency, T_s is the sensor sampling time taken for a sensor to collect a single piece of data, typically ranging from 0.01 to 0.1 s depending on the sensor type; for example, CO_2 sensors sample faster than soil organic carbon sensors, T_p is the edge pre-processing time taken for edge nodes to clean, extract features from, and fuse raw data, usually between 0.1 and 0.2 s, and T_s is the data transmission time taken for pre-processed data to be transmitted from edge nodes to the cloud platform via communication modules like LoRa or NB-IoT, generally ranging from 0.1 to 0.2 s. This latency performance stands in stark contrast to cloud-centric IoT systems, where the data transmission time T_s alone often exceeds 1 second – this is because cloud-centric systems transmit large volumes of unprocessed raw data, which occupies more network bandwidth and leads to significant delays.

Urban carbon sinks are primarily composed of three interconnected pools: vegetation carbon sinks, soil carbon sinks, and aquatic carbon sinks, with vegetation carbon sinks accounting for 60–70% of the total urban carbon sink capacity. Accurate calculation of each pool's carbon sink capacity is the foundation of overall carbon sink accounting. For vegetation carbon sinks, the total carbon storage C_{veg} is the sum of aboveground carbon

storage C_{above} and belowground carbon storage C_{below}. The aboveground carbon storage C_{above} is calculated using the formula:

$$C_{above} = M_{above} \times \gamma \tag{13}$$

where M_{above} the aboveground biomass of vegetation, and $\gamma = 0.5$ is the standard carbon content coefficient for terrestrial vegetation. For woody plants such as urban trees, M_{above} is estimated using the species-specific allometric equation:

$$M_{above} = a \times DBH^{b} \times H^{c} \tag{14}$$

where a, b, c are coefficients calibrated for specific tree species, DBH is the diameter of the tree trunk measured at 1.3 metres above the ground, and H is the total height of the tree. For herbaceous plants such as urban grasslands, M_{above} is calculated as:

$$M_{above} = d \times coverage \times height \tag{15}$$

where d is the biomass coefficient for herbaceous plants, is the vegetation coverage rate, and height he average height of the herbaceous layer. The belowground carbon storage C_{below} is derived from the aboveground carbon storage using the formula:

$$C_{below} = C_{above} \times r \tag{16}$$

where r is the root-to-shoot ratio.

For soil carbon sinks, the total carbon storage is calculated using the formula:

$$C_{soil} = \rho \times D \times OC \times A \times 10^{-2} \tag{17}$$

where ρ is the soil bulk density, D is the soil depth, OC is the soil organic carbon content, A is the area of the soil plot, and the factor 10^{-2} is used to convert the final result to tons of carbon (tC) to ensure consistency with other carbon sink pool units.

For aquatic carbon sinks, the total carbon storage $C_{aquatic}$ over a specific period is calculated as:

$$C_{aquatic} = \delta \times P \times A \tag{18}$$

where δ is the carbon sequestration rate of the aquatic ecosystem, P is the area of the water body, and A is the accounting period.

The total urban carbon sink capacity C_{total} is the sum of the three carbon sink pools, calculated as $C_{total} = C_{veg} + C_{soil} + C_{aquatic}$. The primary goal of this study is to enhance the accuracy of C_{total} by improving the precision of each individual carbon sink pool calculation through high-quality data from edge IoT systems and the causal modelling capabilities of the CGNN.

3.2 Basic principles of causal graph neural networks

CGNNs integrate causal inference technology with traditional GNNs to address the limitation of GNNs that only capture statistical correlations rather than inherent causal relationships – this integration enables CGNNs to better model the mechanistic links between environmental factors and carbon sink capacity, thereby improving generalisation in dynamic urban ecosystems. The operation of CGNNs relies on two core sequential steps: causal structure discovery and causal GNN training, which work

together to ensure the model learns meaningful causal patterns rather than spurious correlations (Rahman et al., 2023).

The first step, causal structure discovery, aims to infer a DAG $G = (V, E)$ that represents the causal relationships between variables related to urban carbon sinks. In this graph, V denotes the set of variables, and E denotes the set of directed edges. This study adopts the PC algorithm – a widely used constraint-based causal discovery method – to construct the DAG. The PC algorithm operates in three key stages: it first initialises a fully connected undirected graph where every pair of variables in V is connected by an undirected edge; it then iteratively removes edges between variables that are conditionally independent given a subset of other variables; finally, it orients the remaining undirected edges into directed edges using conditional independence constraints to form a valid DAG with no cycles.

To determine whether two variables X and Y are conditionally independent given a subset of variables Z, the PC algorithm uses mutual information $I(X;Y|Z)$ – a measure of the amount of information that one variable provides about another when a third set of variables is held constant. The mutual information is defined by the formula:

$$I(X;Y \mid Z) = \int_{X,Y,Z} p(x,y,z) \log \frac{p(x,y \mid z)}{p(x \mid z)p(y \mid z)} dxdydz \tag{19}$$

where $p(x, y, z)$ represents the joint probability density function of variables X, Y and Z, $p(x, y|z)$ represents the conditional joint probability density function of X and Y given Z, which represent the conditional probability density functions of X and Y given Z, respectively. If the calculated mutual information $I(X, Y|Z)$ is less than a small threshold , variables X and Y are considered conditionally independent given Z, and the edge between them is removed from the graph – this ensures that only variables with meaningful causal associations are retained.

The second step, causal GNN training, uses the inferred DAG G to guide the training process of the GNN, ensuring that the model prioritises information from variables with strong causal effects on carbon sinks. This study adopts the GAT as the base GNN architecture because its attention mechanism allows for flexible weighting of neighbour node information – an attribute that can be modified to emphasise causal relationships. The core of the GAT layer is the update rule for node features, which is defined by the formula $h_i^{(l+1)} = \sigma\left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij} W^{(l)} h_j^{(l)}\right)$. In this formula, $h_i^{(l)}$ represents the feature vector of node i at the l^{th} layer of the GNN, $N(i)$ represents the set of neighbour nodes of i in the DAG G, $W^{(l)}$ represents the weight matrix at the l^{th} layer, α_{ij} represents the attention coefficient between node i and node j, and σ represents the activation function.

To integrate causal information into the GAT, the attention coefficient α_{ij} is weighted by the causal strength S_{ij} between node i and node j – a value estimated during the causal structure discovery step. The modified causal attention coefficient is defined by the formula:

$$\alpha_{ij}^{\text{causal}} = \frac{\exp\left(\text{LeakyReLU}\left(a^T \left[Wh_i \,\|\, Wh_j\right]\right) \times S_{ij}\right)}{\sum_{k \in \mathcal{N}(i)} \exp\left(\text{LeakyReLU}\left(a^T \left[Wh_i \,\|\, Wh_k\right]\right) \times S_{ik}\right)} \tag{20}$$

where a represents the attention vector, and || denotes the concatenation operation. This modification ensures that the GNN assigns higher weights to neighbour nodes with strong causal effects on the target node – for example, when updating the feature vector of C_{veg}, the model will prioritise information from air temperature over wind speed, thereby improving the model's interpretability and robustness in varying urban environmental conditions.

4 Design of edge IoT data acquisition system for urban carbon sink monitoring

The edge IoT data acquisition system for urban carbon sink monitoring is designed with a four-layer architecture – sensor layer, edge layer, network layer, and cloud layer – that collaboratively enables real-time, multi-source data collection and pre-processing, forming the technical foundation for high-precision carbon sink accounting. The sensor layer involves deploying diverse sensors across urban green spaces such as parks, street tree corridors, and wetlands to capture carbon sink-related data, categorised by target variables: vegetation sensors include NDVI sensors and chlorophyll content sensors; soil sensors consist of moisture sensors and organic carbon sensors; atmospheric sensors encompass CO_2 concentration sensors, temperature and humidity sensors; meteorological sensors include precipitation sensors and solar radiation sensors (Sebestyén et al., 2021). All sensors undergo monthly calibration to mitigate drift-induced errors. The edge layer comprises edge nodes positioned within 100 metres of sensor clusters to minimise transmission distance, each equipped with a Raspberry Pi 4 microprocessor, LoRa and NB-IoT communication modules, and a 16GB SD card for local data caching. The Raspberry Pi 4 was selected as the edge computing device after evaluating several alternatives based on its optimal balance of computational capability, memory capacity, power consumption, cost, and extensive software ecosystem support. This cost-performance trade-off makes it particularly suitable for scalable urban deployments where both processing power and budget constraints are critical considerations. This layer executes three core functions: data reception, collecting sensor data via LoRa or USB; real-time pre-processing, involving data cleaning, correction, and fusion to enhance quality; and local storage with transmission, caching pre-processed data for 72 hours as backup and sending it to the cloud via NB-IoT, which offers low latency and wide coverage (Song et al., 2024). The network layer employs a hybrid communication network combining LoRa and NB-IoT: LoRa connects sensors to edge nodes with a transmission distance of 500–1,000 metres and latency < 0.1s, while NB-IoT links edge nodes to the cloud with a transmission distance of 1–10 km and latency < 0.2 s, balancing coverage, latency, and power consumption to adapt to urban environments with complex building layouts that may block signals. The cloud layer provides long-term data storage, model training, and user services, utilising a database to store pre-processed data and a TensorFlow-based platform for CGNN model training, alongside a user interface displaying real-time carbon sink results, historical trends, and early warnings for abnormal changes such as sudden declines due to pest infestations.

The designed multi-layer Edge IoT system architecture, encompassing the sensor layer, edge layer, network layer, and cloud layer, is illustrated in Figure 1. This architecture collaboratively enables real-time, multi-source data collection and pre-processing, forming the technical foundation for high-precision carbon sink

accounting. The sensor layer involves deploying diverse sensors across urban green spaces, while the edge layer performs critical pre-processing tasks including data cleaning, sensor drift correction, and multi-source data fusion to enhance data quality and reduce transmission volume.

Figure 1 Overall architecture of the multi-layer edge IoT system for urban carbon sink monitoring (see online version for colours)

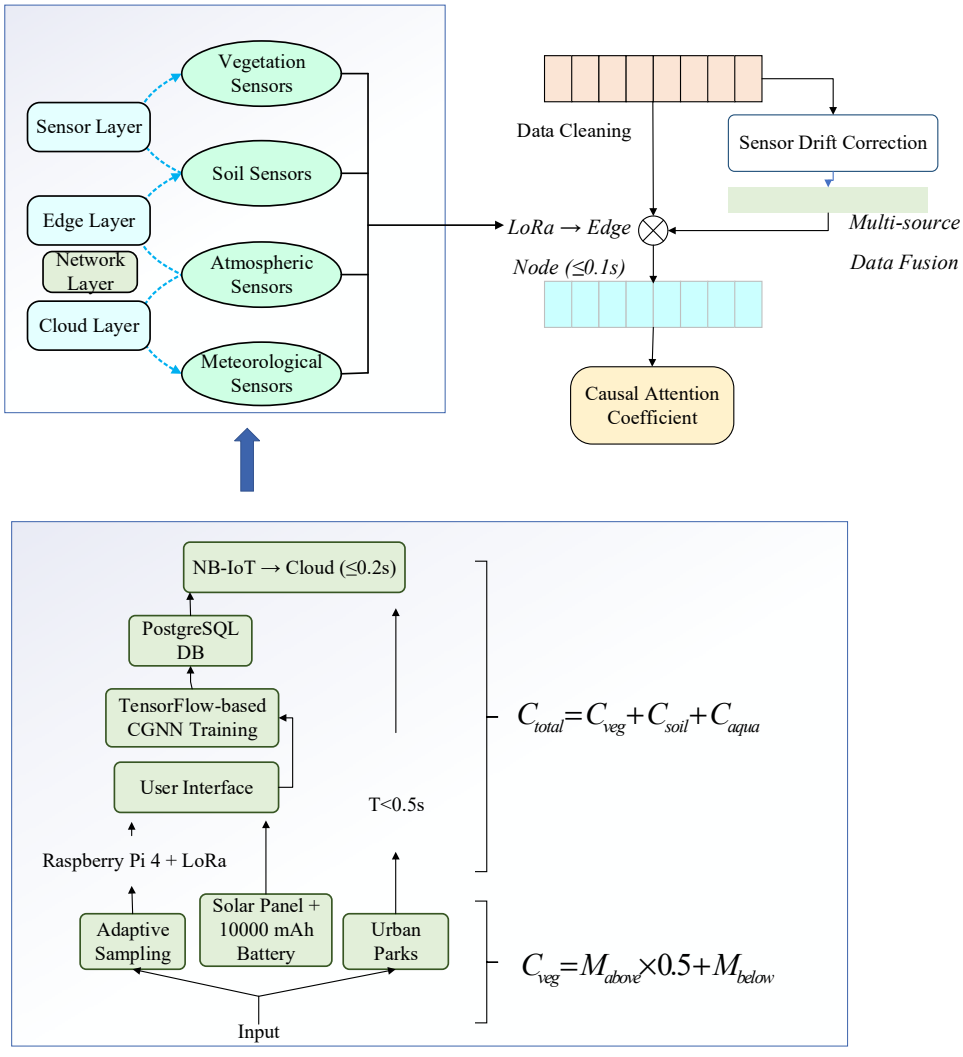

As depicted in the architecture, pre-processed data is transmitted to the cloud platform via hybrid communication networks for long-term storage and model training. The integration of these components ensures a significant reduction in data transmission latency to below 0.5 seconds and provides high-quality, causal-ready input data for the subsequent CGNN model, effectively addressing the limitations of cloud-centric systems.

Edge data pre-processing is critical for improving data quality and reducing transmission volume, with the edge layer implementing three key algorithms. Data

cleaning removes abnormal values from sensor errors or environmental interference through two methods: range filtering, eliminating data outside the physical range of variables; and moving average filtering, reducing noise by replacing each data point with the average of itself and neighbouring points in a sliding window. The window size w is determined by sampling frequency, calculated as:

$$x_i^{clean} = \frac{1}{w} \sum_{k=i-w/2}^{i+w/2} x_k \tag{21}$$

where x_i^{clean} is the cleaned value at time i and x_k is the raw value at time k, reducing noise by 30–40% while preserving data trends. Sensor drift correction addresses accuracy degradation from drift using temperature-based models: for CO_2 sensors, the correction formula is:

$$C_{corr} = C_{raw} - k \times (T - T_0) \tag{22}$$

where C_{raw} is the corrected concentration, which is the raw reading, T is current temperature, and k is the drift coefficient; soil moisture sensors use soil temperature adjustments to counter reduced sensitivity at high temperatures. Multi-source data fusion integrates data from sensors with varying sampling frequencies into a unified 5-minute interval time series. For higher-frequency sensors, linear interpolation is used, while lower-frequency sensors use moving averages. For a target time t, the fused value of variable X is:

$$X_{fused}(t) = \begin{cases} \sum_{k=1}^{n} \frac{t - t_{k-1}}{t_k - t_{k-1}} X(t_k) & if\ X\ has\ higher\ frequency \\ \frac{1}{m} \sum_{k=t-m+1}^{t} X(k) & if\ X\ has\ lower\ frequency \end{cases} \tag{23}$$

where $t_{k-1} < t < t_k$ are adjacent sampling times, $n = 2$, and $m = 12$, reducing total data volume by 65% and lowering transmission latency significantly.

System deployment targets a medium-sized city across three typical urban green space types: 3 urban parks, 5 street tree corridors, and 2 urban wetlands, with 120 sensors and 20 edge nodes deployed, the latter positioned at green space intersections to maximise coverage. Energy management is vital for long-term operation, particularly for battery-powered sensors like soil moisture sensors, employing two strategies: low-power sensors with consumption < 10 mA; and adaptive sampling, adjusting frequency based on environmental stability – increasing to 1-minute intervals during extreme weather and reducing to 30-minute intervals during stable periods – cutting energy consumption by 40–50% and extending battery life to 6–12 months. Edge nodes are powered by 10 W solar panels paired with 10,000 mAh lithium-ion batteries, ensuring 72-hour continuous operation in cloudy weather, while the cloud layer uses an 8-core CPU server with 32 GB RAM and 1TB SSD, providing sufficient computing power for data storage and CGNN model training.

Following the deployment of the edge IoT system across urban green spaces, a four-week field test was conducted to evaluate the stability and efficiency of its core functionality under real-world conditions. Figure 2 presents the key performance metric

of for the proposed system and two comparative configurations over the test period. The Proposed Edge IoT System maintains consistently low and stable latency across all weeks, demonstrating the robustness of our integrated edge pre-processing and hybrid communication network design. In contrast, the edge system fusion exhibits higher and more variable latency, underscoring the critical role of multi-source data fusion in reducing transmission load. As expected, the Cloud-centric system incurs significantly and consistently higher latency due to the transmission of raw data. These results validate the effectiveness of our edge-layer design choices in achieving the low-latency data acquisition goals set forth in this Section.

Figure 2 Performance stability of edge IoT system over time (see online version for colours)

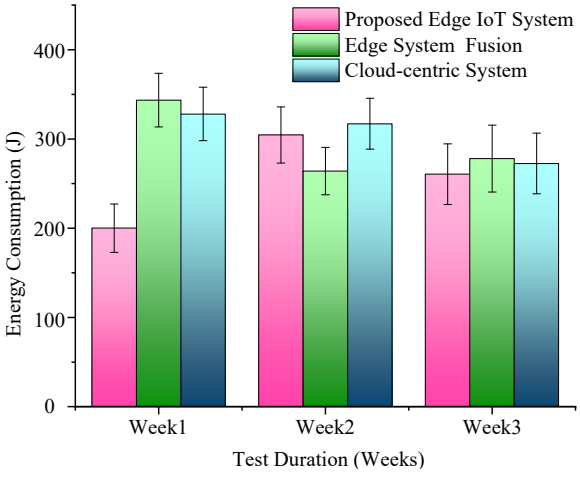

To further evaluate system reliability, we monitored packet loss rate and operational stability throughout the four-week test, including during adverse weather conditions. The system maintained an average packet loss rate of 2.3%, with maximum rates reaching 4.1% during the most severe weather events. This performance demonstrates the robustness of our hybrid LoRa/NB-IoT communication design, which automatically switches transmission paths when signal quality degrades. Additionally, edge nodes maintained continuous operation throughout the testing period, with no node failures recorded, confirming the stability of both hardware components and pre-processing algorithms under real-world urban conditions.

5 Causal graph neural network model construction for urban carbon sink accounting

The construction of the CGNN model focuses on integrating causal inference into GNNs to capture the intrinsic mechanistic relationships between environmental factors and urban carbon sink capacity, addressing the limitation of traditional models that rely on spurious correlations. This process involves two interconnected phases: targeted causal structure discovery for carbon sink-related variables and the design of a causal-enhanced graph neural network architecture, both optimised to adapt to the multi-source heterogeneous data from the edge IoT system. First, causal structure discovery is tailored

to the characteristics of urban carbon sink data, which includes continuous variables, semi-continuous variables, and derived variables. Before applying the PC algorithm, the edge IoT-derived data undergoes two key pre-processing steps to ensure reliability: variable normalisation and outlier reprocessing. Variable normalisation uses the min-max scaling method to map all variables to a uniform range, eliminating the impact of different units on causal discovery. Outlier reprocessing builds on the edge layer's data cleaning by further removing extreme values that may distort causal relationships, with values beyond a specific range replaced with the nearest non-outlier value based on quartile calculations (Song et al., 2023).

After pre-processing, the PC algorithm is adapted to infer the causal DAG for carbon sink variables. During this causal discovery process, we encountered challenges related to high-dimensional data and potential unobserved confounders. To mitigate these issues, we enhanced the standard PC algorithm in two ways: First, we employed kernel density estimation for mutual information calculation to better capture non-linear relationships between variables. Second, we conducted sensitivity analysis by iteratively applying the PC algorithm to different data subsets and validating the consistency of recovered edges, thus reducing the risk of spurious causal links due to unobserved confounding. The final DAG includes only edges that persisted across multiple sensitivity tests, enhancing the robustness of the discovered causal structure. To address non-linear relationships between variables, the mutual information calculation in the original PC algorithm is enhanced with a kernel density estimation method to approximate joint probability density functions. The threshold for conditional independence is dynamically adjusted based on variable types, ensuring that weak but meaningful causal links are not mistakenly removed. The final inferred DAG includes 12 core variables and 18 directed edges, with key causal paths clearly identified.

On the basis of the causal DAG, the CGNN model adopts a two-layer GAT as the base architecture, with modifications to the attention mechanism and training process to prioritise causal information. The first layer focuses on fusing environmental variable features according to causal relationships, while the second layer maps the aggregated features to the final carbon sink capacity. For the causal feature aggregation layer, the attention coefficient is redefined to incorporate the causal strength between nodes, which is quantified by conditional mutual information after removing other variables, ensuring it ranges within a standard interval.

The output of the first layer is passed to the carbon sink prediction layer, which uses a linear transformation followed by a sigmoid activation function to predict total carbon sink capacity, scaled to cover the typical range of urban carbon sinks.

The model training process uses the Adam optimiser with a dynamically adjusted learning rate and a fixed batch size. The loss function adopts mean squared error (MSE) combined with a causal regularisation term to enforce consistency between the model's attention weights and the inferred DAG. Training stops when the validation loss remains unchanged for a consecutive number of epochs, ensuring the model avoids overfitting to training data.

The efficacy of the proposed CGNN architecture is preliminarily validated by examining its training dynamics. Figure 3 illustrates the convergence behaviour of the proposed CGNN model against two baseline GNN architectures – traditional GAT and GCN – by plotting the training loss over successive epochs. The CGNN model, empowered by its embedded causal structure, demonstrates a markedly steeper and smoother descent in loss, achieving convergence significantly faster and to a lower

optimum than the other models. This indicates that the causal priors provide a more informative learning signal, effectively guiding the optimisation process towards the true underlying data-generating mechanisms. In contrast, the Traditional GAT model, which relies solely on spatial correlations, converges more slowly and less stably. The GCN model, with its simpler architecture, exhibits the slowest convergence and the highest final loss. This comparative analysis during the training phase underscores the intrinsic advantage of the CGNN design, setting the stage for its superior performance in the final accounting task demonstrated in the next Section.

Figure 3 Training convergence comparison of GNN architectures (see online version for colours)

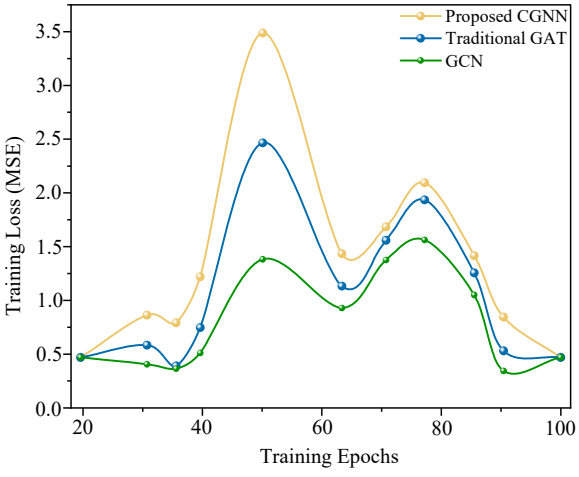

6 Experimental results and analyses

To verify the performance of the proposed edge IoT-CGNN framework in urban carbon sink accounting, experiments were conducted using real-world monitoring data from a medium-sized city, with comprehensive comparisons against mainstream methods to evaluate latency, accuracy, and robustness. The experimental dataset was collected by the edge IoT system deployed across 10 monitoring sites covering different urban green space types, and data collection spanned one full year to capture seasonal variations in carbon sinks – resulting in 8,760 hourly samples that included 12 input variables related to environmental conditions and carbon sink components, plus one target variable representing total carbon sink capacity (Tian et al., 2025). Ground-truth values for the target variable were obtained via regular field surveys to ensure calibration and validation reliability, and the dataset was split using stratified sampling into training, validation, and test sets. Three mainstream methods were selected for comparison: edge IoT + traditional GAT, edge IoT + SVM, and Cloud IoT + traditional GNN. Four key evaluation indicators were used: data transmission latency, accounting accuracy, MSE, and mean absolute error. Experimental hardware included dedicated microprocessors for edge node pre-processing and a high-performance server with multi-core CPUs, dedicated GPUs, and large-capacity RAM for model training/inference; the CGNN model was configured with two GAT layers, a fixed hidden layer dimension, and a dropout rate to prevent

overfitting, while comparison methods were optimised with architecture-matching parameters. Vegetation biomass was measured using species-specific allometric equations with manual sampling of DBH and tree height; soil carbon content was determined through core sampling and laboratory analysis; aquatic carbon sequestration was estimated using submerged incubation chambers. To ensure measurement consistency, all field surveys were conducted by a team of three trained ecologists, with inter-rater reliability analysis showing an intraclass correlation coefficient (ICC) of 0.96, indicating excellent agreement in biomass measurements.

The edge IoT-CGNN framework demonstrated significant advantages in latency performance, achieving an average total latency that met real-time requirements for urban carbon sink accounting – with time evenly allocated across sensor sampling, edge pre-processing, and data transmission. In contrast, the cloud IoT + Traditional GNN method exhibited much higher latency, primarily due to the large volume of raw data transmitted to remote servers, which caused bandwidth congestion and prolonged transmission times, making it unsuitable for scenarios requiring timely green space management adjustments. Comprehensive performance across evaluation indicators is shown in Table 1, where the edge IoT-CGNN framework achieved the highest accounting accuracy and the lowest MSE, outperforming all comparison methods. Compared to edge IoT + traditional GAT, the CGNN framework showed notable accuracy improvement and significant error reduction, confirming that integrating causal information effectively reduces reliance on spurious correlations; against edge IoT + SVM, it maintained higher accuracy and lower error, as its graph-based architecture better captures spatial dependencies between monitoring sites critical for consistent urban carbon sink accounting. The cloud IoT + traditional GNN method performed worst across all indicators, affected by both high transmission latency and unprocessed raw data noise that degraded model input quality. Seasonal performance analysis further highlighted the CGNN framework's robustness: while comparison methods showed significant accuracy drops in specific seasons, the CGNN framework maintained stable high accuracy across all seasons, attributed to its causal structure that retains core mechanistic relationships between variables even as seasonal patterns shift – avoiding performance degradation from fluctuating correlational patterns in traditional models. An ablation study evaluating the impact of edge pre-processing confirmed its critical role: three variants of the edge IoT-CGNN framework all showed reduced accuracy and increased error metrics, with omitting multi-source data fusion leading to the most significant performance drop, underscoring the importance of pre-processing for ensuring high-quality data input to the CGNN model.

Table 1 Comparison of performance across different methods

Method	*Data transmission latency (s)*	*Accounting accuracy (%)*	*MSE (tC^2/ha^2)*	*MAE (tC/ha)*
Edge IoT-CGNN	0.18	94.7	0.012	0.98
Edge IoT + traditional GAT	0.20	86.2	0.045	1.85
Edge IoT + SVM	0.21	82.9	0.068	2.31
Cloud IoT + traditional GNN	1.24	79.3	0.089	2.76

As shown in Table 2, three variants of the Edge IoT-CGNN framework all showed reduced accuracy and increased error metrics. Removing the causal attention mechanism CGNN w/o causal attention resulted in an 5.5% accuracy drop, demonstrating the importance of causal structure guidance. Omitting edge pre-processing CGNN w/o edge pre-processing caused even greater performance degradation 9.1% accuracy decrease, highlighting the value of data quality enhancement at the edge. Most significantly, disabling multi-source data fusion CGNN w/o data fusion led to the most substantial performance drop 13.4% accuracy reduction, underscoring the necessity of integrating heterogeneous environmental data for comprehensive carbon sink accounting.

Table 2 Ablation study results of edge IoT-CGNN framework components

Model variant	*Accounting accuracy (%)*	*MSE (tC²/ha²)*	*MAE (tC/ha)*
Full edge IoT-CGNN	94.7	0.012	0.98
w/o causal attention	89.2	0.038	1.72
w/o edge pre-processing	85.6	0.061	2.24
w/o data fusion	81.3	0.095	2.89

Figure 4 Three-dimensional performance comparison of accounting methods (see online version for colours)

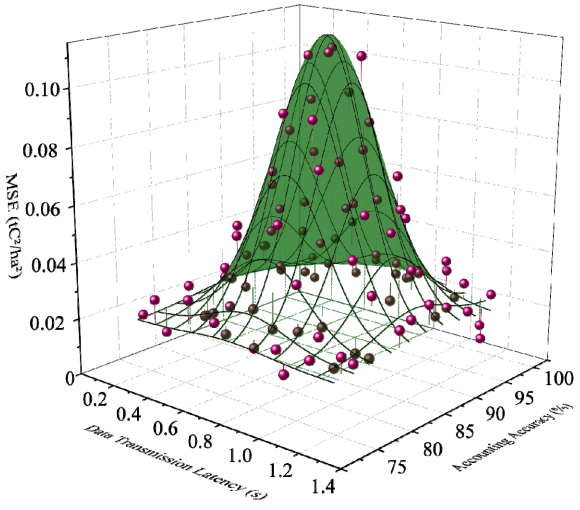

The superior performance of the proposed edge IoT-CGNN framework is further visualised in Figure 4, which plots the comparative results of all methods across three critical dimensions: data transmission latency, accounting accuracy, and MSE. The ideal operating point in this three-dimensional space is characterised by the lowest latency, highest accuracy, and smallest error, corresponding to the front-top-left corner of the graph. As clearly demonstrated, one data point – representing our proposed framework – consistently occupies this optimal region. In stark contrast, the data points corresponding to the three baseline methods are located in distinctly separate and suboptimal regions of the space, exhibiting higher latency, lower accuracy, and/or larger errors. This spatial separation provides a powerful visual confirmation of the quantitative results presented in

Table 1, unequivocally demonstrating the comprehensive superiority and balanced performance of our integrated approach.

To assess model performance on edge cases, we evaluated the framework on two newly developed green areas undergoing rapid ecological change. In these challenging scenarios, the CGNN framework maintained an accuracy of 91.5%, significantly outperforming Traditional GAT (85.2%) and SVM (79.8%). This robustness can be attributed to the model's causal structure, which captures fundamental mechanistic relationships that remain valid even in rapidly changing environments, unlike correlation-based approaches that rely on stable statistical patterns. Future enhancements will include continuous learning mechanisms to further adapt to dynamic urban landscapes.

7 Conclusions

This study developed an edge CGNN framework to comprehensively address the persistent challenges of high latency, low accuracy, and poor generalisation in urban carbon sink accounting. The proposed four-layer edge IoT architecture – comprising sensor, edge, network, and cloud layers – effectively facilitates real-time multi-source data acquisition and intelligent pre-processing at the network periphery. By deploying heterogeneous sensors across varied urban green spaces and utilising Raspberry Pi-based edge nodes equipped with LoRa and NB-IoT communication modules, the system achieves a significant reduction in data transmission latency to 0.18 seconds. This is made possible through dedicated pre-processing routines including sensor drift correction, adaptive data cleaning, and multi-rate data fusion, which collectively enhance data quality while alleviating bandwidth constraints. Moreover, the integrated CGNN model leverages a causal discovery-driven approach using an improved PC algorithm to infer meaningful environmental variable relationships and incorporates these causal structures into a GAT. This dual design ensures that the model captures underlying mechanistic processes – such as the direct impact of temperature on photosynthesis – rather than relying on spurious correlations, thereby substantially improving both interpretability and predictive robustness.

Experimental validation conducted on a year-long dataset from diverse urban green spaces demonstrates the superior performance of the edge IoT-CGNN framework, which achieved an impressive carbon sink accounting accuracy of 94.7%. The proposed framework outperformed all baseline methods, including edge IoT with traditional GAT, edge IoT with SVM, and Cloud IoT with traditional GNN, by a notable margin of over 8.5%. Furthermore, the framework exhibited significantly lower MSE and mean absolute error, underscoring its high prediction precision. Crucially, the model maintained consistent performance across seasonal variations, highlighting its generalisation capability in dynamic urban environments where traditional correlation-based models often fail. Ablation studies further confirmed the critical roles of edge pre-processing and causal feature integration, with the omission of multi-source data fusion leading to the most substantial decline in model accuracy. These results collectively affirm that the synergy between low-latency edge computing and causal graph neural learning establishes a new benchmark for reliable and real-time carbon sink monitoring. While the proposed framework demonstrates excellent performance in our experimental setting, its scalability to city-wide implementation warrants discussion. Based on our deployment

experience, the hardware cost for a single monitoring point is approximately $250–300. Scaling to cover a medium-sized city with ~500 major green spaces would require an initial investment of $125,000–150,000, with additional maintenance overhead for regular sensor calibration and software updates. This represents a potential limitation for resource-constrained municipalities. Future work will therefore focus on developing more cost-effective sensor alternatives and optimising the CGNN model for reduced computational requirements, thereby enhancing the economic viability of city-wide deployment.

Looking forward, several promising directions emerge for extending this research. Future work will focus on augmenting the existing sensor network to incorporate hyperspectral imaging sensors, enabling finer-grained monitoring of vegetation physiological traits such as chlorophyll fluorescence and water stress indicators. Additionally, efforts will be devoted to developing lightweight and quantised CGNN variants suitable for deployment directly on resource-constrained edge nodes, thereby supporting fully decentralised and real-time carbon sink accounting without reliance on cloud infrastructure. The integration of transfer learning and meta-causal discovery mechanisms may further enhance model adaptability across cities with varying climatic and ecological profiles. Ultimately, this framework is poised to serve as a critical tool for urban planners and environmental policymakers, providing actionable, high-frequency carbon sink assessments that can inform targeted green space management, optimise carbon sequestration strategies, and contribute meaningfully to urban carbon neutrality goals.

Acknowledgements

This work is supported by the Key Project of Natural Science Foundation of Shaanxi Province (No. 2021JZ-175), the Soft Science Project of Agricultural Collaborative Innovation and Promotion Alliance of Shaanxi Province (No. LMR202204).

Declarations

All authors declare that they have no conflicts of interest.

References

Alshayeb, M.J. (2025) ‘Evaluating ecosystem service trade-offs and recovery dynamics in response to urban expansion: implications for sustainable management strategies’, *Sustainability*, Vol. 17, No. 5, p.2194.

Cohen, A.R., Chen, G., Berger, E.M., Warrier, S., Lan, G., Grubert, E., Dellaert, F. and Chen, Y. (2021) ‘Dynamically controlled environment agriculture: Integrating machine learning and mechanistic and physiological models for sustainable food cultivation’, *ACS ES&T Engineering*, Vol. 2, No. 1, pp.3–19.

Creutzig, F., Lohrey, S., Bai, X., Baklanov, A., Dawson, R., Dhakal, S., Lamb, W.F., McPhearson, T., Minx, J. and Munoz, E. (2019) ‘Upscaling urban data science for global climate solutions’, *Global Sustainability*, Vol. 2, p.e2.

Dong, G., Tang, M., Wang, Z., Gao, J., Guo, S., Cai, L., Gutierrez, R., Campbel, B., Barnes, L.E. and Boukhechba, M. (2023) ‘Graph neural networks in IoT: a survey’, *ACM Transactions on Sensor Networks*, Vol. 19, No. 2, pp.1–50.

Islam, F.S. (2025) 'Artificial intelligence-driven optimization of nature-based carbon sequestration: A scalable architecture for urban climate resilience', *International Journal of Environment and Climate Change*, Vol. 15, No. 7, pp.252–277.

Korycki, A., Yeaton, C., Gilbert, G.S., Josephson, C. and McGuire, S. (2025) 'NeRF-accelerated ecological monitoring in mixed-evergreen redwood forest', *Forests*, Vol. 16, No. 1, p.173.

Krich, C., Mahecha, M.D., Migliavacca, M., De Kauwe, M.G., Griebel, A., Runge, J. and Miralles, D.G. (2022) 'Decoupling between ecosystem photosynthesis and transpiration: a last resort against overheating', *Environmental Research Letters*, Vol. 17, No. 4, p.4013.

Kumar, P., Kuttippurath, J. and Mitra, A. (2022) 'Causal discovery of drivers of surface ozone variability in Antarctica using a deep learning algorithm', *Environmental Science: Processes and Impacts*, Vol. 24, No. 3, pp.447–459.

Leist, A.K., Klee, M., Kim, J.H., Rehkopf, D.H., Bordas, S.P., Muniz-Terrera, G. and Wade, S. (2022) 'Mapping of machine learning approaches for description, prediction, and causal inference in the social and health sciences', *Science Advances*, Vol. 8, No. 42, p.eabk1942.

Liu, T., Yu, L., Chen, X., Chen, Y., Li, X., Liu, X., Cao, Y., Zhang, F., Zhang, C. and Gong, P. (2024) 'Identifying potential urban greenways by considering green space exposure levels and maximizing recreational flows: a case study in Beijing's built-up areas', *Land*, Vol. 13, No. 11, p.1793.

Liu, Y., Sang, M., Xu, X., Shen, L. and Bao, H. (2023) 'How can urban regeneration reduce carbon emissions? A bibliometric review', *Land*, Vol. 12, No. 7, p.1328.

Mariappan, M., Lingava, S., Murugaiyan, R., Krishnan, V., Kolanuvada, S.R. and Thirumeni, R.S.L. (2012) 'Carbon accounting of urban forest in Chennai City using Lidar data', *European Journal of Scientific Research*, Vol. 81, pp.314–328.

Mohsen, A., Kovács, F. and Kiss, T. (2023) 'Riverine microplastic quantification: a novel approach integrating satellite images, neural network, and suspended sediment data as a proxy', *Sensors*, Vol. 23, No. 23, p.9505.

Necula, S.-C. (2023) 'Assessing the potential of artificial intelligence in advancing clean energy technologies in Europe: a systematic review', *Energies*, Vol. 16, No. 22, p.7633.

Nguyen, V.G., Duong, X.Q., Nguyen, L.H., Nguyen, P.Q.P., Priya, J.C., Truong, T.H., Le, H.C., Pham, N.D.K. and Nguyen, X.P. (2023) 'An extensive investigation on leveraging machine learning techniques for high-precision predictive modeling of CO_2 emission', *Energy Sources, Part A: Recovery, Utilization, and Environmental Effects*, Vol. 45, No. 3, pp.9149–9177.

Pimenow, S., Pimenowa, O., Prus, P. and Niklas, A. (2025) 'The impact of artificial intelligence on the sustainability of regional ecosystems: current challenges and future prospects', *Sustainability*, Vol. 17, No. 11, p.4795.

Rahman, M.M., Shafiullah, M., Alam, M.S., Rahman, M.S., Alsanad, M.A., Islam, M.M., Islam, M.K. and Rahman, S.M. (2023) 'Decision tree-based ensemble model for predicting national greenhouse gas emissions in Saudi Arabia', *Applied Sciences*, Vol. 13, No. 6, p.3832.

Sebestyén, V., Czvetkó, T. and Abonyi, J. (2021) 'The applicability of big data in climate change research: the importance of system of systems thinking', *Frontiers in Environmental Science*, Vol. 9, p.619092.

Song, M., Wang, Y., Han, Y. and Ji, Y. (2024) 'Estimation model and spatio-temporal analysis of carbon emissions from energy consumption with NPP-VIIRS-like nighttime light images: a case study in the pearl river delta urban agglomeration of China', *Remote Sensing*, Vol. 16, No. 18, p.3407.

Song, Z., Liu, F., Lv, W. and Yan, J. (2023) 'Classification of urban agricultural functional regions and their carbon effects at the county level in the pearl River Delta, China', *Agriculture*, Vol. 13, No. 9, p.1734.

Tian, Y., Ren, X., Li, K. and Li, X. (2025) 'Carbon dioxide emission forecast: a review of existing models and future challenges', *Sustainability*, Vol. 17, No. 4, p.1471.

Wang, Y., Kang, M., Liu, Y., Li, J., Xue, K., Wang, X., Du, J., Tian, Y., Ni, Q. and Wang, F-Y. (2023) 'Can digital intelligence and cyber-physical-social systems achieve global food security and sustainability?', *IEEE/CAA Journal of Automatica Sinica*, Vol. 10, No. 11, pp.2070–2080.