



**International Journal of Global Warming**

ISSN online: 1758-2091 - ISSN print: 1758-2083

<https://www.inderscience.com/ijgw>

---

**Topology-aware modelling of regional energy systems using hierarchical graph networks for carbon peaking analysis**

Yun Zhang

**DOI:** [10.1504/IJGW.2026.10080329](https://doi.org/10.1504/IJGW.2026.10080329)

**Article History:**

|                   |                 |
|-------------------|-----------------|
| Received:         | 21 January 2026 |
| Last revised:     | 27 June 2026    |
| Accepted:         | 30 June 2026    |
| Published online: | 13 August 2026  |

---

# Topology-aware modelling of regional energy systems using hierarchical graph networks for carbon peaking analysis

---

Yun Zhang

School of Economics and Management,  
Zhengzhou Normal University,  
Zhengzhou, Henan, 450044, China  
Email: zhangy2026zznu@outlook.com

**Abstract:** A physics-constrained hierarchical graph network is developed for carbon forecasting based on directed energy-flow graphs. The model combines topology-specific attention with temporal encoding, hierarchical pooling and an energy-balance penalty to capture interregional flow dependence and thermodynamic consistency. Using 30-province data from 2000–2022, it achieves a root mean square error of 14.32 million tonnes of CO<sub>2</sub>, 27.86% lower than the strongest baseline, a coefficient of determination of 0.98, and an energy imbalance rate of 0.42%. Attribution analysis assigns 42.00% of emission growth to coal-to-power flows. Under Shared Socioeconomic Pathway 2, Jiangsu emissions decline after 2020, whereas Inner Mongolia peaks in 2029.

**Keywords:** graph neural networks; regional energy systems; energy-flow graph; carbon peaking; physics-constrained learning; dynamic attribution.

**Reference** to this paper should be made as follows: Zhang, Y. (2026) 'Topology-aware modelling of regional energy systems using hierarchical graph networks for carbon peaking analysis', *Int. J. Global Warming*, Vol. 39, No. 5, pp.1–10.

**Biographical notes:** Yun Zhang graduated from the Bournemouth University, UK in 2013. She studied in tourism and hospitality management, Bournemouth University. Her research interests include management.

---

## 1 Introduction

China has committed to peaking carbon dioxide emissions before 2030 and achieving carbon neutrality before 2060 (Government of the People's Republic of China, 2021). Achieving this target requires provincial mitigation decisions within an interconnected energy system in which coal, oil, gas, electricity, and heat are transferred across administrative boundaries before final use. Provincial emissions are therefore shaped by local demand, upstream conversion, interprovincial transfer, and imported energy. A directed energy-flow graph provides a suitable representation for this system because it preserves the physical sequence from resource supply to transformation and final consumption (Cullen and Allwood, 2010).

Regional carbon forecasting has developed from econometric analysis toward nonlinear sequence learning and graph-based modelling. Regression models based on population, affluence, and technology provide transparent macro-level baselines (York et al., 2003). Long short-term memory networks capture sequential dependence in energy and emission trajectories (Hochreiter and Schmidhuber, 1997). Transformers further improve long-range temporal modelling through self-attention (Vaswani et al., 2017). Recent studies have introduced relational learning into carbon and energy-system prediction. Pan et al. (2024) used improved graph convolution for day-ahead carbon-emission prediction in power systems. Zhao et al. (2025) combined graph convolution with temporal convolution for real-time grid-emission estimation. Shao et al. (2026) developed a graph-based multi-branch framework for urban carbon mapping. These studies confirm the value of relational structure, but grid connectivity and spatial adjacency do not fully represent the resource, conversion, transfer, and consumption relationships contained in provincial energy-balance sheets.

Two limitations remain in existing forecasting frameworks. Forecasting error alone cannot determine whether predicted supply, transformation, and final consumption satisfy energy-balance constraints. Physics-informed learning addresses this issue by incorporating known physical residuals into the optimisation objective (Raissi et al., 2019). Interpretability is also limited. The logarithmic mean Divisia index, abbreviated as LMDI, provides residue-free decomposition through a predefined accounting identity (Ang, 2005), but it is primarily designed for historical attribution. Integrated gradients can attribute the output of a differentiable model to input features relative to a selected baseline (Sundararajan et al., 2017), which provides a potential basis for dynamic attribution in nonlinear graph models.

This study proposes a physics-constrained hierarchical graph network, abbreviated as PC-HGN, for regional energy consumption forecasting and carbon-peaking analysis. The central novelty lies in aligning graph topology with provincial energy-balance accounting rather than using geographic adjacency or generic grid connectivity alone. The model combines type-aware graph encoding, temporal aggregation, hierarchical pooling, and differentiable penalties for conversion and system-balance residuals. A model-attribution module termed neural-LMDI aggregates node- and edge-level integrated gradients into scale, structure, and intensity effects. The objectives are to evaluate out-of-sample forecasting accuracy against statistical, sequence, and graph baselines, quantify physical consistency and component contributions, and analyse regional heterogeneity and conditional carbon-peaking trajectories.

## 2 Methodology

The framework converts provincial energy statistics into directed heterogeneous graphs, encodes their spatial and temporal evolution, and applies energy-balance penalties during training. Model outputs are then examined through additive gradient attribution.

### 2.1 Energy-flow graph construction

For province  $r$  and year  $t$ , the provincial energy balance sheet is represented as a directed energy-flow graph. Accounting categories for production, transformation, transfer, final

consumption, and loss follow the Energy Statistics Manual (International Energy Agency, 2005). Equation (1) defines the graph.

$$G_{r,t} = (V_R \cup V_T \cup V_C, E_{r,t}, A_{r,t}, X_{r,t}) \quad (1)$$


The resource, transformation, and final-consumption node sets are mutually exclusive. The edge set records observed forward transfers, the weighted adjacency matrix stores their magnitudes, and the feature matrix stores node attributes. Admissible links include resource-to-transformation, resource-to-consumption, transformation-to-transformation, and transformation-to-consumption flows. Transfer-in and transfer-out entries are retained explicitly so that each provincial graph remains connected to interregional exchanges.

Each node combines lagged endogenous states, node-specific exogenous drivers, and province-level context, as expressed in equation (2).

$$x_{i,r,t} = h_{i,r,t}^{endo} \parallel h_{i,r,t}^{exo} \parallel u_{r,t} \quad (2)$$

The left-hand side of equation (2) is the node feature vector. Endogenous states include lagged energy flows and emissions. Exogenous drivers are selected by node type and include resource availability, conversion efficiency, installed capacity, gross domestic product, population, and sector-composition variables. Gross domestic product is abbreviated as GDP. The global context contains province-level variables shared across nodes. All energy quantities are converted to a common accounting unit before graph construction.

**Figure 1** Energy-flow graph construction (see online version for colours)



## 2.2 Hierarchical graph and temporal encoding

The spatial encoder follows the heterogeneous graph transformer of Hu et al. (2020). Separate projections are used for resource, transformation, and consumption nodes, while relation-specific parameters distinguish conversion, direct-use, and transfer edges. Equation (3) gives the attention coefficient for head  $h$ .

$$\alpha_{ij}^h = \text{softmax}_j \left( \frac{(q_i^h)^T (k_j^h + \tilde{e}_{ij}^h)}{\sqrt{d_h}} \right) \quad (3)$$

The numerator combines the type-specific query, key, and projected edge representations, and the denominator scales the score by the attention-head dimension. Equation (4) aggregates neighbouring node and edge information. Residual feed-forward processing and layer normalisation then produce the spatial representation for each year.

$$H_{r,t}^{macro} = D^{-1} S BiGRU(Z_{r,t-L+1:t}) \quad (4)$$

Temporal evolution is modelled with a bidirectional gated recurrent unit, abbreviated as BiGRU, following the gated recurrent formulation of Cho et al. (2014). For each forecast, the encoder is applied only to the historical input window preceding the target year. A fixed assignment matrix maps detailed nodes to aggregate sectors, and a degree matrix provides normalisation. Equation (5) gives the resulting macro-scale representation.

$$\mathcal{L}_{conv} = |V_T|^{-1} \sum_{v \in V_T} (\hat{y}_{v,t}^{out} - \eta_v \hat{y}_{v,t}^{in})^2 \quad (5)$$

Parallel decoders generate node-level flows, aggregate-sector consumption, and provincial carbon emissions. Joint prediction across these scales links detailed energy pathways to the provincial output used in carbon-peaking analysis.

### 2.3 Physical regularisation and neural-LMDI

Physical consistency is introduced through differentiable penalty terms. Equation (6) penalises conversion imbalance at transformation nodes.

$$\mathcal{L}_{conv} = |V_T|^{-1} \sum_{v \in V_T} (\hat{y}_{v,t}^{out} - \eta_v \hat{y}_{v,t}^{in})^2 \quad (6)$$

The predicted output of each transformation node is compared with its predicted input multiplied by the observed or specified conversion efficiency. Equation (7) penalises imbalance among net supply, terminal consumption, and modelled loss.

$$\mathcal{L}_{sys} = [\hat{S}_t - (\hat{C}_t + \hat{L}_t)]^2 \quad (7)$$

The complete objective is given by Equation (8).

$$\mathcal{L} = \mathcal{L}_{pred} + \lambda(e) (\mathcal{L}_{conv} + \mathcal{L}_{sys}) + \gamma \|\Theta\|_2^2 \quad (8)$$

The prediction term combines node-, sector-, and province-level errors. The epoch-dependent coefficient  $\lambda(e)$  increases gradually during training,  $\gamma$  controls L2 parameter regularisation, and  $\Theta$  denotes the trainable parameter set. The schedule gives the data term greater influence at the beginning of training and progressively increases the contribution of the physical penalties.

Neural-LMDI attributes the change in predicted emissions between consecutive years. The previous-year feature state is used as the baseline, and integrated gradients is

evaluated along the straight path to the current state. Equation (9) defines the attribution for feature  $k$ .

$$IG_k = (x_k - x_{k'}) \int_0^1 \frac{\partial F[x' + \alpha(x - x')]}{\partial x_k} d\alpha \quad (9)$$

In equation (9),  $F$  denotes the trained PC-HGN,  $x$  is the current state,  $x'$  is the baseline, and alpha indexes the integration path. Feature-level attributions are aggregated into GDP, population, sectoral structure, energy intensity, and edge-flow effects. The completeness property supports additive accounting of the model-output difference. Neural-LMDI should therefore be interpreted as a nonlinear attribution procedure rather than as the logarithmic-mean decomposition identity used in classical LMDI.

### 3 Experimental setup

The evaluation uses a chronological test design and common inputs across competing models. Predictive accuracy, physical consistency, robustness, architectural contribution, regional heterogeneity, attribution, and scenario behaviour are examined separately.

#### 3.1 Data and evaluation

The dataset covers 30 provincial-level administrative divisions from 2000 to 2022. Tibet, Hong Kong, Macao, and Taiwan are excluded because complete annual balance sheets are unavailable in the selected statistical series. Energy-flow and socioeconomic observations are obtained from the China energy and statistical yearbooks issued from 2001 to 2023 (National Bureau of Statistics of China, 2001–2023). The panel contains 690 province-year graphs. Each graph has approximately 45 heterogeneous nodes, including 18 energy commodities and the principal transformation and consumption categories, with approximately 250 nonzero directed edges.

Carbon emissions are calculated from terminal fuel consumption using the 2006 Guidelines for National Greenhouse Gas Inventories issued by the Intergovernmental Panel on Climate Change (2006). Observations from 2018 to 2022 form the held-out test period, and earlier observations are used for model development. Root mean square error, mean absolute error, and the coefficient of determination evaluate predictive accuracy. The metrics are denoted RMSE, MAE, and  $R^2$ , respectively. Errors are reported in million tonnes of carbon dioxide, denoted Mt CO<sub>2</sub>. Physical consistency is evaluated with the energy imbalance rate, abbreviated as EIR, defined in equation (10).

$$EIR = \frac{100}{T} \sum_{t=1}^T \frac{|\hat{S}_t - (\hat{C}_t + \hat{L}_t)|}{\hat{S}_t} \quad (10)$$

In equation (10),  $T$  is the number of test observations. Lower EIR values indicate closer agreement with the modelled supply-demand balance. All reported decimal values use two digits after the decimal point.

### 3.2 Baselines and ablation design

Six baselines cover statistical, econometric, sequence, and graph forecasting. Autoregressive integrated moving average, abbreviated as ARIMA, follows Box et al. (2015). The stochastic impacts by regression on population, affluence, and technology model, abbreviated as STIRPAT, follows York et al. (2003) and uses ridge regularisation. The sequence baselines are long short-term memory, abbreviated as LSTM, and the transformer (Hochreiter and Schmidhuber, 1997; Vaswani et al., 2017). The graph convolutional network with LSTM, abbreviated as GCN-LSTM, combines graph convolution (Kipf and Welling, 2017) with recurrent encoding. The spatio-temporal graph convolutional network, abbreviated as STGCN, follows Yu et al. (2018). All trainable models use the same input variables and chronological split.

Three ablation variants isolate the main model components. The first removes the two physical penalties, the second replaces heterogeneous encoding with homogeneous graph convolution, and the third removes hierarchical pooling. These variants distinguish the effects of physical regularisation, relation-aware encoding, and multiscale aggregation.

## 4 Results and analysis

Results are reported for the 2018 to 2022 held-out period. Statistical, sequence, and graph baselines are compared first, followed by physical validation, robustness, regional analysis, attribution, and conditional scenario projections.

### 4.1 Predictive accuracy and physical consistency

Figure 2 reports the aggregated results across the 30 provinces. PC-HGN obtains an RMSE of 14.32 Mt CO<sub>2</sub>, an MAE of 10.13 Mt CO<sub>2</sub>, and a coefficient of determination of 0.98. STGCN is the lowest-error baseline, with an RMSE of 19.85 Mt CO<sub>2</sub>, an MAE of 14.22 Mt CO<sub>2</sub>, and a coefficient of determination of 0.95. The corresponding reductions are 27.86% for RMSE and 28.62% for MAE. These results indicate that the energy-flow topology contains predictive information beyond temporal covariates alone.

**Figure 2** (a) RMSE and MAE (b) Coefficient of determination (see online version for colours)





Figure 3(a) shows an EIR of 0.42% for PC-HGN, compared with 4.98% for STGCN, 5.43% for GCN-LSTM, 6.82% for LSTM, and 7.15% for the Transformer. The result is consistent with the intended function of the balance penalties. Figure 3(b) separates the component effects. Removing the physical-regularisation terms raises RMSE from 14.32 to 15.10 Mt CO<sub>2</sub> and EIR from 0.42% to 4.85%. Replacing heterogeneous encoding produces the largest accuracy loss, with RMSE increasing to 16.80 Mt CO<sub>2</sub>. Removing hierarchical pooling increases RMSE to 15.50 Mt CO<sub>2</sub>. The physical term contributes mainly to balance consistency, while relation-aware encoding has the largest effect on predictive error.

**Figure 3** (a) Energy imbalance rate (b) Ablation results (see online version for colours)



## 4.2 Robustness and regional heterogeneity

The 2020 observations provide a distribution-shift test because transport and industrial activity changed abruptly. Relative to the corresponding pre-2020 error, the LSTM error increase exceeds 40.00%, whereas the increase for PC-HGN remains below 15.00%. The graph representation allows changes in terminal demand to propagate to upstream resource and transformation nodes, while the physical penalties discourage isolated responses that violate the modelled balance.


Regional performance is examined after provinces are grouped by the sign of net interprovincial energy transfer in the balance sheet. In energy-exporting provinces, the RMSE reduction relative to the lowest-error baseline exceeds 30.00%. These provinces contain longer extraction, processing, conversion, and transfer chains, so explicit topology provides a stronger structural signal. The improvement is smaller in energy-importing coastal provinces, where terminal demand accounts for a larger share of the graph. Imported electricity remains represented as an external resource flow, preserving the associated upstream dependence.

## 4.3 Dynamic attribution and scenario analysis

Figure 4(a) presents the Neural-LMDI attribution for Shanxi from 2015 to 2020. GDP and population contribute 28.50 and 5.20 Mt CO<sub>2</sub>, while sectoral structure and energy intensity contribute −12.40 and −18.60 Mt CO<sub>2</sub>. The edge-flow effect contributes 22.10 Mt CO<sub>2</sub>, giving a net increase of 24.80 Mt CO<sub>2</sub>. The result indicates that a lower


industrial share can coexist with stronger dependence on coal-based electricity. The attribution analysis in Figure 4(b) assigns 42.00% of national emission growth to the coal-to-power flow, followed by electricity-to-steel at 15.00% and coke-to-steel at 12.00%. These values identify the energy paths with the largest model-attributed contributions.

**Figure 4** (a) Shanxi driver attribution (b) National energy-path attribution (see online version for colours)



Scenario inputs follow the qualitative ordering of the shared socioeconomic pathways described by O'Neill et al. (2017). Figure 5 reports shared socioeconomic pathway 2, abbreviated as SSP2. The Jiangsu trajectory declines from its 2020 starting value, whereas Inner Mongolia reaches a modelled maximum in 2029. Under the uniform policy setting, western energy bases require a modelled annual coal-output adjustment of  $-8.00\%$ . The differentiated setting reduces the adjustment to  $-2.50\%$ . The comparison indicates a substantially steeper transition requirement under a common deadline. These values are conditional model outputs under the stated scenario assumptions and are not deterministic forecasts.

**Figure 5** (a) Provincial trajectories under SSP2 (b) Annual coal-output adjustment (see online version for colours)



## 5 Conclusions

PC-HGN represents provincial energy systems as directed heterogeneous graphs and combines graph, temporal, hierarchical, and balance-regularised learning. During the 2018 to 2022 test period, it obtains an RMSE of 14.32 Mt CO<sub>2</sub>, an MAE of 10.15 Mt CO<sub>2</sub>, a coefficient of determination of 0.98, and an EIR of 0.42%. Removing the physical penalty increases EIR to 4.85%, while homogeneous graph encoding increases RMSE to 16.80 Mt CO<sub>2</sub>. The attribution analysis links 42.00% of emission growth to coal-to-power flows. Under SSP2, Jiangsu emissions decline after 2020 and Inner Mongolia reaches a modelled maximum in 2029. The results support topology-aware forecasting and region-specific peaking assessment within the limits of the available data and scenario assumptions.

## Declarations

The authors declare no conflicts of interest.

## References

- Ang, B.W. (2005) 'The LMDI approach to decomposition analysis: a practical guide', *Energy Policy*, Vol. 33, No. 7, pp.867–871, <https://doi.org/10.1016/j.enpol.2003.10.010>.
- Box, G.E.P., Jenkins, G.M., Reinsel, G.C. and Ljung, G.M. (2015) *Time Series Analysis: Forecasting and Control*, 5th ed., Wiley, Hoboken, New Jersey.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H. and Bengio, Y. (2014) 'Learning phrase representations using RNN encoder-decoder for statistical machine translation', in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pp.1724–1734, Association for Computational Linguistics, <https://doi.org/10.3115/v1/D14-1179>
- Cullen, J.M. and Allwood, J.M. (2010) 'The efficient use of energy: tracing the global flow of energy from fuel to service', *Energy Policy*, Vol. 38, No. 1, pp.75–81, <https://doi.org/10.1016/j.enpol.2009.08.054>.
- Government of the People's Republic of China (2021) *China's Achievements, New Goals and New Measures for Nationally Determined Contributions*, United Nations Framework Convention on Climate Change [online]. Available at: <https://unfccc.int/documents/307765> (accessed 20 January 2026).
- Hochreiter, S. and Schmidhuber, J. (1997) 'Long short-term memory', *Neural Computation*, Vol. 9, No. 8, pp.1735–1780, <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Hu, Z., Dong, Y., Wang, K. and Sun, Y. (2020) 'Heterogeneous graph transformer', in *Proceedings of The Web Conference 2020*, pp.2704–2710, Association for Computing Machinery, <https://doi.org/10.1145/3366423.3380027>.
- Intergovernmental Panel on Climate Change (2006) *2006 IPCC Guidelines for National Greenhouse Gas Inventories*, in Eggleston, S., Buendia, L., Miwa, K., Ngara, T. and Tanabe, K. (Eds.), Institute for Global Environmental Strategies [online] <https://www.ipcc-nggip.iges.or.jp/public/2006gl/> (accessed 20 January 2026).
- International Energy Agency (2005) *Energy Statistics Manual*, OECD Publishing, <https://doi.org/10.1787/9789264033986-en>.
- Kipf, T.N. and Welling, M. (2017) 'Semi-supervised classification with graph convolutional networks', in *International Conference on Learning Representations* [online] <https://openreview.net/forum?id=SJU4ayYGL> (accessed 20 January 2026).

- National Bureau of Statistics of China (2001–2023) *China Statistical Yearbook and China Energy Statistical Yearbook [Annual Statistical Series]*, China Statistics Press, Beijing.
- O'Neill, B.C., Krieglner, E., Ebi, K.L., Kemp-Benedict, E., Riahi, K., Rothman, D.S., van Ruijven, B.J., van Vuuren, D.P., Birkmann, J., Kok, K., Levy, M. and Solecki, W. (2017) 'The roads ahead: narratives for shared socioeconomic pathways describing world futures in the 21st century', *Global Environmental Change*, Vol. 42, pp.169–180, <https://doi.org/10.1016/j.gloenvcha.2015.01.004>.
- Pan, F., Yang, Y., Ji, Y., Li, J., Zhang, J. and Zhong, L. (2024) 'Application of improved graph convolutional networks in daily-ahead carbon emission prediction', *Frontiers in Energy Research*, Vol. 12, p.1371507, <https://doi.org/10.3389/fenrg.2024.1371507>.
- Raissi, M., Perdikaris, P. and Karniadakis, G.E. (2019) 'Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations', *Journal of Computational Physics*, Vol. 378, pp.686–707, <https://doi.org/10.1016/j.jcp.2018.10.045>.
- Shao, W., Tu, Y., Yue, T., Zhao, N., Zhou, H. and Zhang, L. (2026) 'Mapping Chinese urban carbon emissions with a graph-based multi-branch deep learning framework', *GIScience & Remote Sensing*, Vol. 63, No. 1, p.2631837, <https://doi.org/10.1080/15481603.2026.2631837>.
- Sundararajan, M., Taly, A. and Yan, Q. (2017) 'Axiomatic attribution for deep networks', in *Proceedings of the 34th International Conference on Machine Learning*, Vol. 70, pp.3319–3328, PMLR [online] <https://proceedings.mlr.press/v70/sundararajan17a.html> (accessed 20 January 2026).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł. and Polosukhin, I. (2017) 'Attention is all you need', in *Advances in Neural Information Processing Systems*, Vol. 30, pp.5998–6008, Curran Associates.
- York, R., Rosa, E.A. and Dietz, T. (2003) 'STIRPAT, IPAT and ImPACT: analytic tools for unpacking the driving forces of environmental impacts', *Ecological Economics*, Vol. 46, No. 3, pp.351–365, [https://doi.org/10.1016/S0921-8009\(03\)00188-5](https://doi.org/10.1016/S0921-8009(03)00188-5).
- Yu, B., Yin, H. and Zhu, Z. (2018) 'Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting', in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, International Joint Conferences on Artificial Intelligence Organization*, pp.3634–3640, <https://doi.org/10.24963/ijcai.2018/505>.
- Zhao, Q., Chen, J., Jia, Q., Sun, C., Chen, X. and Chen, H. (2025) 'Predicting real-time carbon emissions for power grids using graph convolutional networks', *Machines*, Vol. 13, No. 11, p.1061, <https://doi.org/10.3390/machines13111061>.