

Improving spatial allocation for energy system coupling with graph neural networks^{*}

Xuanhao Mu^{ID}^{*}, Jakob Geiges^{ID}, Nan Liu^{ID}, Thorsten Schlachter^{ID}, Veit Hagenmeyer^{ID}

Institute for Automation and Applied Informatics, Karlsruhe Institute of Technology, Germany

ARTICLE INFO

Keywords:

Deep learning
Energy system coupling
Granularity gap
Graph neural network
Spatial resolution

ABSTRACT

In energy system analysis, coupling models with mismatched spatial resolutions is a significant challenge. A common solution is assigning weights to high-resolution geographic units for aggregation, but traditional models are limited by using only a single geospatial attribute. This paper presents an innovative method employing a self-supervised Heterogeneous Graph Neural Network to address this issue. This method models high-resolution geographic units as graph nodes, integrating various geographical features to generate physically meaningful weights for each grid point. These weights enhance the conventional Voronoi-based allocation method, allowing it to go beyond simply geographic proximity by incorporating essential geographic information. In addition, the self-supervised learning paradigm overcomes the lack of accurate ground-truth data. Experimental results demonstrate that applying weights generated by this method to cluster-based Voronoi Diagrams significantly enhances scalability, accuracy, and physical plausibility, while increasing precision compared to traditional methods. Code is available at <https://github.com/KIT-IAI/AllocateGNN>

1. Introduction

In the context of global climate change and increasing energy demand, the energy transition and decarbonization have become critical issues. Accurate and comprehensive energy system simulation models are essential tools for addressing these challenges. The models frequently exhibit varying spatial resolutions, so a key issue in model coupling is the efficient conversion of one model's outputs at a specific spatial granularity into inputs suitable for other models operating at different spatial granularities. Unlike the traditional spatial disaggregation problem, its ultimate goal is not to generate a high resolution raster map, but to allocate the total amount to a set of nodes or areas with a specific distribution.

Mu et al. [1] used **Grid Point Modeling (GPM)** with **Clustering Induced Voronoi Diagram (CIVD)** to address the spatial load allocation process, which initially disaggregates the total volume into high resolution intermediate units and subsequently allocates these units to the final target entities utilizing techniques such as the **Voronoi Diagram (VD)**. Although this method is superior to traditional **VD**, it still has the following core limitations in its allocation mechanism:

- 1) Grid point weights are calculated only based on their primary land use type. The grid points of the same type (e.g., all residential points) are given the same weight.
- 2) **GPM** is a static process that struggles to effectively integrate multiple land use proportions or other multi source heterogeneous information.

The main contributions of this paper are: First, we propose a **GPM** framework based on heterogeneous **Graph Neural Network (GNN)**, which integrates multiple geospatial features to generate data-driven allocation weights. This framework exhibits good scalability and paves the way for the future introduction of more auxiliary data. Second, we design a self-supervised learning paradigm that uses macroeconomic and socioeconomic indicators as supervisory signals, addressing the key challenge of lacking ground-truth data in the real world. Finally, experiments demonstrate that combining this framework with the **CIVD** method significantly improves spatial allocation accuracy, reducing the **Root Mean Squared Error (RMSE)** by an average of 4.87 %. Code is available at <https://github.com/KIT-IAI/AllocateGNN>

This paper is structured as follows: **Section 2** presents a review of related work. **Section 3** outlines the methods used in detail. **Section 4**

^{*} The Helmholtz Association supported this work under the project Helmholtz platform for the design of robust energy systems and their supply chains" (RESUR).

^{*} Corresponding author.

E-mail addresses: xuanhao.mu@kit.edu (X. Mu), jakob.geiges@kit.edu (J. Geiges), nan.liu@kit.edu (N. Liu), thorsten.schlachter@kit.edu (T. Schlachter), veit.hagenmeyer@kit.edu (V. Hagenmeyer).

<https://doi.org/10.1016/j.epsr.2026.113519>

Received 2 October 2025; Received in revised form 10 February 2026; Accepted 20 June 2026

Available online 30 June 2026

0378-7796/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

describes the data set used in this paper and evaluates the results of the presented methods. Section 5 concludes a discussion, and Section 6 provides a conclusion and the future work.

2. Related work

Energy system simulation models play a crucial role in addressing the dual challenges of global climate change and growing energy demand. However, with the increasing prevalence of Distributed Energy Resources (DERs) and the increasing complexity of energy networks, these models generally face granularity gaps [2]. Spatial resolution is a core dimension of this gap, directly impacting the reliability and practicality of model results [3]. Improving spatial resolution is crucial for optimizing power system costs and planning [4], but it must be balanced against computational feasibility [5].

The allocation of macro-regional totals to micro-geographic units is crucial for bridging data across different models. A classic method is spatial disaggregation, which means breaking spatially aggregated data into smaller spatial units to reveal local features. Common spatial disaggregation methods include Areal Weighting [6], Dasymetric Mapping [7], Machine Learning [8], and geostatistical methods, such as Kriging Interpolation [9].

Nonetheless, these traditional spatial disaggregation techniques frequently prove insufficient for directly addressing the granularity gap in model coupling. Their main goal is usually to produce a continuous high-resolution raster map, instead of allocating a total quantity to a defined set of discrete nodes or regions needed by an alternative model. Hülk et al. [10] use VD to address this specific allocation challenge. They allocate the annual German electricity consumption and production to the corresponding substations using a two-stage process. First, using land use data, the total regional volume is spatially disaggregated into several high-spatial-resolution load areas. Secondly, the load data is allocated to the corresponding substations.

A VD divides the space into regions, where each region contains all the locations closest to a given point. Clustering Induced Voronoi Diagram (CIVD) [11] is a novel Voronoi Diagram method that improves the traditional space allocation by dividing spaces using clusters. The value of k nearest points for each region in the VD is calculated by clustering. However, these geometric-based methods inherently rely on pre-defined heuristics (such as distance and proximity) and are unable to dynamically and adaptively learn the complex drivers of spatial allocation.

To incorporate more socioeconomic factors, the method proposed by Hülk et al. [10] utilizes land use data to create a weighted distribution. Mu et al. [1] used GPM with CIVD to improve this spatial load allocation process. The GPM method is first used to calculate the weight distribution based on land use for each grid point, which is obtained by dividing the entire area to be allocated into small squares and has a higher spatial resolution compared to the load area mentioned in [10]. Then, CIVD assigns these grid points to substations through clustering and distance-based rules. Lastly, the final allocation result (e.g., load distribution) is calculated based on the total weight allocated. While this method considers more factors than solely geometric methods, its assigned weights are typically calculated statically based on predefined rules. This complicates the capacity of the model to capture the complex nonlinear relationships between different geographic features. To overcome these limitations, Graph Neural Networks (GNNs) have been introduced as a powerful framework for learning complex spatial relationships.

GNNs are specifically designed to process graph-structured data and can automatically learn complex patterns and relationships from node features and graph topology. In recent years, GNNs have demonstrated significant potential in a variety of geospatial problems, including urban planning [12], traffic flow prediction [13], etc.

By modeling geographic regions and power grid nodes as a heterogeneous graph, GNNs are able to learn an end-to-end mapping function. Unlike traditional models like Graph Convolutional Networks (GCNs) [14], which can only process homogeneous nodes, heteroge-

neous GNNs, such as the Heterogeneous Graph Transformer (HGT) [15], can effectively handle different node types (e.g., administrative districts, substations, and commercial centers) and the rich relationships between them. This enables the model to dynamically generate optimal allocation weights for each micro-unit based on multi-source heterogeneous geographical and economic characteristics.

3. Method

The present paper proposes a Graph Neural Network (GNN)-based Grid Point Method (GPM) to generate allocation weights for grid points in geographic space. The core idea is to model the spatial allocation problem as a graph learning task. The overall pipeline consists of four core components: Graph Builder, Graph Encoder, Edge Weight Predictor, and Loss Function, as shown in Fig. 1.

3.1. Graph builder

As the task involves two functionally distinct spatial entities, macro-regions (source nodes) and micro-geographic units (agent nodes), with an asymmetric allocation relationship, we construct a heterogeneous graph to explicitly model this relationship. The graph consists of two different types of nodes:

- Source Nodes $\mathcal{N}(s)$: These represent the macro-geographic regions whose total volume (e.g., the region's annual total electricity consumption) needs to be decomposed and allocated. Node features are primarily derived from macro-statistics of the region, such as total population and Gross Value Added (GVA) of various industries, to describe the socioeconomic profile of the region.
- Agent Nodes $\mathcal{N}(a)$: These represent the micro-geographic units that receive the allocated volume. Uniform, fixed size grids are generated across the entire macro-geographic region, using the center of each grid as an agent node. This discretization method transforms the continuous geographic space into a manageable set of nodes. Node features are primarily derived from geographic information data, such as OpenStreetMap (OSM). Specifically, including the area percentage




Fig. 1. Flowchart of the GNN-based spatial allocation method. The main pipeline (blue) generates weights, while the self-supervised loop (red) computes the loss for training by reconstructing macro-indicators.

of different land use types (e.g., residential and industrial) within the grid, as well as a One-Hot encoding of the dominant land use type.

The connections between nodes, named edges, are crucial to defining how information flows within the graph. In the present paper, edges are constructed from direct geographic affiliations. If the geographic coordinates of an agent node are within the polygonal boundary of a source node, a connection is established between them. To enable bidirectional information exchange, we establish two directed edges for each pair of connected nodes: one from the source node to the agent node, representing the flow of macro-region information to micro-units; and another from the agent node back to the source node, allowing the model to aggregate the states of micro-units to update its understanding of the macro-region.

3.2. Graph encoder

The core task of the Graph Encoder is to learn an embedding vector for each node in the graph. This task faces two major challenges: the original features of the source and agent nodes are heterogeneous, with different dimensions, types, and semantic spaces; the complex relationships defined by different node types and edges in the graph need to be effectively captured. To address these challenges, the encoder adopts a two-stage architecture. First, separate linear layers project the original features of source and agent nodes into a unified latent space. Subsequently, we use a multi-layer **Heterogeneous Graph Transformer (HGT)** as the core information aggregation module. **HGT** is an attention architecture explicitly designed for heterogeneous graphs and can learn unique semantic weights for each relationship based on the types of nodes and edges.

Finally, the encoder generates an embedding vector for each node that combines its own characteristics with graph structure information, providing high quality input for downstream weight prediction tasks.

3.3. Edge weight predictor

After the Graph Encoder generates node embeddings, the Edge Weight Predictor converts these abstract representations into concrete assigned weights. This fully differentiable prediction module aims to predict a weight w_{sa} for each edge connecting a source node $s \in \mathcal{N}(s)$ with an agent node $a \in \mathcal{N}(a)$, ensuring that the sum of the weights of all edges originating from the same source node is strictly 1, i.e. $\sum_{a \in \mathcal{N}(s)} w_{sa} = 1$. This module includes Relation Cost Calculation and Grouped Softmax Normalization.

In the embedding space, an agent node should be semantically closer to the source node to which it belongs. This relationship cost $c(s, a)$ can be quantified by a gated distance. It is obtained by multiplying two parts:

- **Euclidean Distance:** Directly computes the distance $d(s, a)$ between the source node embedding h_s and the agent node embedding h_a .
- **Gating Score:** A small **Multi-Layer Perceptron (MLP)** takes the concatenated vector of the two node embeddings $[h_s || h_a]$ as input and outputs a relationship score $g(s, a)$. This score serves as a learnable gating signal to adjust the importance of distance dynamically.

The final relationship cost is calculated as follows:

$$c(s, a) = \underbrace{\sigma(\text{MLP}([h_s || h_a]))}_{g(s, a): \text{Gating Score}} \cdot \underbrace{\|h_s - h_a\|_2}_{d(s, a): \text{Euclidean Distance}} \quad (1)$$

where $[h_s || h_a]$ represents vector concatenation, and σ is a sigmoid function that ensures the gating score is between (0, 1). Lower cost values $c(s, a)$ indicate a better match between agent node a and source node s .

In order to convert the “lower is better” cost value into a “higher is better” weight that “sums to 1”, it is necessary to apply the Softmax function to the negative costs of all outgoing edges of each source node

s:

$$w_{sa} = \frac{\exp(-c(s, a)/\tau)}{\sum_{a' \in \mathcal{N}(s)} \exp(-c(s, a')/\tau)} \quad (2)$$

where $a' \in \mathcal{N}(s)$ is the set of all agent nodes connected to the source node s . τ is a temperature hyperparameter that controls the smoothness of the weight distribution.

Through the above design, the entire edge weight predictor is fully differentiable. This ensures that the gradient of the loss can be smoothly back-propagated to each layer of the graph encoder.

3.4. Loss function

One of the primary challenges in spatial allocation is the lack of ground-truth data. This means that for consumption allocation problems, although some aggregated load data may exist at the substation level, the precise and true load values for each individual micro-geographic unit are usually unavailable. To overcome this bottleneck, a self-supervised learning framework is used. Instead of relying on direct labels, this framework generates supervisory signals by defining agent tasks based on the intrinsic structure of the data and domain knowledge, thereby indirectly but effectively guiding model training.

The core loss function is the land use prediction loss. This loss function stems from a core domain assumption: the **GVA** and population share of different functional areas (e.g., industrial and commercial areas) within a macro-region should be highly correlated with the distribution of energy consumption across these functional categories [10]. Therefore, if the distribution weights w_{sa} learned by the model are meaningful, it should be able to reconstruct the macro-distribution of socioeconomic indicators by aggregating the attributes of agent nodes $\mathcal{N}(a)$.

First, each source node s constructs a true, normalized distribution vector P_s .

$$P_s = \left[\frac{\text{GVA}_{\text{industry}}}{\text{Total}}, \frac{\text{GVA}_{\text{commerce}}}{\text{Total}}, \dots, \frac{\text{Population}}{\text{Total}} \right]_s \quad (3)$$

Then, the model uses its predicted edge weights w_{sa} to perform a weighted summation of the land use type of each agent node a (represented as a One-Hot vector T_a) to reconstruct the distribution $\hat{P}_s$.

$$\hat{P}_s = \sum_{a \in \mathcal{N}(s)} w_{sa} T_a \quad (4)$$

Finally, **Kullback-Leibler divergence (KL-divergence)** is used to measure the difference between the reconstructed distribution and the true distribution P . The total loss is the sum of **KL-divergence** over all source nodes:

$$\mathcal{L}_{\text{self-supervised}} = \sum_{s \in S} D_{KL}(P_s || \hat{P}_s) \quad (5)$$

3.5. Clustering introduced Voronoi diagram

The **GNN** model proposed in this paper serves as a core weight generation module and is integrated with the **CIVD** method [1]. This constitutes a two-stage spatial allocation process, enabling allocation from macro-regional aggregate demand to specific power facilities. The weighted agent nodes generated by **GNN** will serve as input to the **CIVD** allocation process. The **CIVD** method first clusters power facilities (e.g., substations) in geographic space to form multiple load centers. Then, for each agent node, the algorithm assigns it to the nearest load center based on its Euclidean distance to each load center. Finally, the weighted agent nodes within each cluster are aggregated and evenly distributed to power facilities within that region.

For the implementation and theoretical basis of the **CIVD** method, please refer to research [1] and [11].

Table 1
Comparison of root mean squared error (RMSE) performance across various geographic regions.

Region	VD	VD-GPM	VD-GNN-GPM	CIVD	CIVD-GPM	CIVD-GNN-GPM
Training Set						
London	21.085	16.234	17.013 (−4.80%)	19.933	16.220	15.998 (+1.37%)
TLH2	8.805	6.735	6.656 (+1.17%)	7.647	5.816	5.778 (+0.65%)
TLH3	9.269	7.265	7.064 (+2.77%)	9.069	6.727	6.535 (+2.85%)
TLJ1	15.831	9.981	9.979 (+0.02%)	13.676	8.917	8.493 (+4.76%)
TLF1	11.966	8.591	8.129 (+5.38%)	10.840	7.627	7.072 (+7.28%)
TLF2	12.574	9.301	9.266 (+0.38%)	10.884	7.542	6.970 (+7.58%)
TLC1	18.302	9.965	9.494 (+4.73%)	17.630	8.882	7.873 (+11.36%)
TLC2	13.318	9.936	9.484 (+4.55%)	8.173	5.841	5.820 (+0.36%)
TLD6	8.674	6.139	5.991 (+2.41%)	7.769	5.384	5.017 (+6.82%)
TLG1	15.492	11.041	10.861 (+1.63%)	14.763	10.740	9.983 (+7.05%)
TLG2	13.909	9.563	9.032 (+5.55%)	12.984	11.795	11.694 (+0.86%)
TLE4	13.770	7.449	7.430 (+0.26%)	13.353	7.007	6.688 (+4.55%)
Test Set						
TLH1	7.532	5.514	5.636 (−2.21%)	7.188	5.670	5.407 (+4.64%)
TLE3	17.336	9.357	9.521 (−1.75%)	12.720	7.479	6.870 (+8.14%)
TLD3	10.201	7.068	7.153 (−1.20%)	9.592	6.579	6.266 (+4.76%)
TLD4	11.537	5.636	5.830 (−3.44%)	9.808	4.424	4.443 (−0.43%)

4. Results

The present paper uses the Great Britain primary substations dataset from [16] to verify the effectiveness of the described method. This dataset contains the peak demand at 11 kV substations in various main regions of Great Britain. The experiments used data from 16 different International Territorial Level (ITL) [17] regions. The train set included 12 regions: London (TLI3-7), TLH2 (Cheshire), TLH3 (Greater Manchester), TLJ1 (West Midlands), TLF1 (Derbyshire and Nottinghamshire), TLF2 (Leicestershire, Rutland and Northamptonshire), TLC1 (Cumbria), TLC2 (Lancashire), TLD6 (South Yorkshire), TLG1 (Herefordshire, Worcestershire and Warwickshire), TLG2 (Staffordshire and Shropshire) and TLE4 (North Northamptonshire). The test set included the last 4 regions not included in training: TLH1 (Merseyside), TLE3 (Lincolnshire), TLD3 (South Yorkshire), and TLD4 (West Yorkshire). Due to its ultra-dense nature, London is evaluated as a separate region, which can be used to test the model's performance in extreme cases. Each region is divided into approximately 50,000 grid points at different grid resolutions. All calculations in this paper were executed using Python. Geographic data was processed and analyzed with the GeoPandas Python library [18]. The Geographic Information System (GIS) data is sourced from OSM [19] and is retrieved using the OSMnx [20] Python library. Population statistics and GVA figures for administrative regions were obtained from sources [21] and [22].

4.1. Allocation performance

Table 1 compares the performance of our proposed method with several baseline methods across various geographic regions, using the RMSE as the evaluation metric. Lower RMSE values indicate superior model performance. The abbreviations for the methods in the table are as follows: VD (Voronoi Diagram), GPM (Grid Point Modeling), and CIVD (Cluster Introduced Voronoi Diagram). To clearly present the comparison results, bold numbers in the table indicate the best performance among all methods in that region. The percentages in parentheses quantify the performance improvement of the GNN for GPM, representing the percentage reduction achieved by the GNN enhancement method (e.g., CIVD-GNN-GPM) relative to its corresponding GPM baseline (e.g., CIVD-GPM). A positive value represents a performance optimization (RMSE reduction), while a negative value indicates performance degradation (RMSE increase, indicated with an underline).

The model demonstrated good learning and fitting capabilities on the training set. The CIVD-GNN-GPM method consistently outperforms the previous CIVD-GPM [1] method without the GNN across all train-




Fig. 2. Land use distribution with substation in region TLD4.

ing regions, achieving an average RMSE reduction of 4.87%. For example, in the TLC1 region, the application of GNN-GPM reduced the RMSE from 8.882 to 7.873, corresponding to an error reduction of 11.36 %. This steady improvement is significant because the GNN model doesn't use any real substation data as a supervisory signal. Furthermore, the training was based solely on open-source data from OSM. It demonstrates that the model can successfully learn physically meaningful and realistic load spatial distribution patterns based solely on the distribution patterns of macroeconomic and socioeconomic indicators. This improvement is not only a statistical optimization but also represents a significant improvement in the fidelity of the load distribution model. This also lays a good foundation for the future introduction of additional auxiliary data.

However, the performance of the overall two-stage allocation framework hinges on the compatibility of the methodologies between modules. This is exemplified by comparing GNN-enhanced methods with traditional Voronoi Diagram (VD) baselines. For example, the CIVD-GNN method almost always outperforms CIVD. However, in some areas, VD-GPM underperforms VD. For example, in London, the application of GNN-GPM improves the RMSE from 17.013 to 16.234. This




Fig. 3. Land use distribution and weights heat map for region London (top) and TLC1 (Bottom).

phenomenon can be attributed to a synergistic mismatch between the weight generation and spatial allocation stages of the two-stage allocation method. GPM can generate a weight distribution that is close to the actual demand. However, traditional VD strictly follows geometric partitioning, and its rigid partition boundaries may cause spatial misalignment with the high-load areas identified by GNN. This misalignment can lead to error amplification: a high-load area identified by GNN may be completely misattributed to a suboptimal facility by the VD boundary, resulting in a larger error than the fuzzy average allocation. This comparison proves that the overall performance of the hybrid model depends not only on the advancement of individual modules, but more importantly on the compatibility of the methodologies between modules.

This compatibility issue is also reflected in the test set. The CIVD-GNN-GPM approach maintains its performance advantage when applied to four geographic regions not included in the training set. Specifically, this approach achieves lower RMSE values than CIVD-GPM across three test regions, except TLD4. Numerically, TLD4's performance difference is 0.019, representing a 0.43 % increase in RMSE. Fig. 2 illustrates that the upper part of TLD4 (shown by the green circle) comprises a substantial "other" region devoid of a substation at its center. This region probably utilizes no electricity; nonetheless, both the GNN model and the GPM approach attribute weights to it. This is because neither method knows the exact location of the substation. However, GNNs, by learning the different attributes of grid points, produce a sharper weight distribution than GPM. This sharp weight distribution, combined with the rigid distribution, amplifies the error, leading to performance degradation. It highlights the need to improve the compatibility of hybrid models in the absence of real, labeled data.

4.2. Physical plausibility

To verify the physical plausibility of the proposed method, Fig. 3 qualitatively demonstrates the land use distribution and weight distribution generated by the GNN for two different geographic regions: London

(top) and the TLC1 region (bottom). Using London as an example, the heat map (top right) shows a high concentration of weights in the urban core (green circle), which is consistent with the high density of commercial (yellow) and residential (blue) areas shown in the land use map (top left). Moreover, for the TLC1 region (bottom), which contains large tracts of agricultural land (cyan), the model assigns significantly higher weights to the eastern region (green circle), which is home to a concentration of industrial (red) and residential (blue) areas. In contrast, the weights for the vast agricultural areas and other undeveloped areas are negligible. Overall, these visualizations demonstrate the model's ability to learn the complex relationship between geospatial characteristics and energy demand, resulting in a spatially consistent and physically plausible allocation scheme.

5. Discussion

The results of this study clearly demonstrate that combining GNN with CIVD can significantly improve the accuracy of spatial allocation of energy systems. By abstracting geographic space into a heterogeneous graph, GNNs are able to capture the complex, nonlinear relationships between multidimensional features such as land use, population density, and economic activity, thereby generating dynamic, data-driven allocation weights for each micro-geographic unit. Furthermore, compared to the static GPM approach, GNN can incorporate more auxiliary data through node features to dynamically calculate weights, demonstrating good scalability. Moreover, the model training does not require any ground-truth data and is entirely self-supervised.

While the proposed approach has achieved positive results, it still has some limitations. First, the model's performance is highly dependent on the quality and granularity of the input data, such as the coverage and accuracy of OSM data. Its effectiveness may be limited in data-sparse areas. Subsequently, there are performance differences when GNN is combined with different spatial delineation methods. In the TLD4 region, the GNN, due to its lack of the locations of substations, produced a more concentrated weight distribution in

an anomalous area lacking a substation. In this case, higher accuracy became a disadvantage, amplifying the overall error. This reveals the need to prioritize compatibility between modules when designing a multi-stage framework, rather than simply superimposing advanced technologies.

6. Conclusion

To address the critical spatial granularity gap in coupled energy systems, this paper proposes an innovative two-stage spatial allocation framework based on Graph Neural Network (GNN). The key contribution of this method lies in the first application of heterogeneous GNN to Grid Point Modeling (GPM), which is a method for generating high-resolution spatial distribution weights based on geographic features. A self-supervised learning paradigm is designed to enable this method to generate physically reasonable allocation weights by learning from the distribution patterns of macro-socioeconomic indicators, even in the absence of ground-truth data. Furthermore, the node features in the GNN provide good scalability for the future introduction of additional auxiliary data.

Experimental results validate the effectiveness of this approach. By combining the dynamic weights generated by the GNN with Clustering Induced Voronoi Diagram (CIVD), the proposed CIVD-GNN-GPM model outperforms previous methods in accuracy and demonstrates good generalization on unseen test data. This demonstrates that this framework not only improves allocation accuracy but also provides a new and practical method for solving spatial allocation problems. However, our self-supervised objective relies on macro-proxies to supervise micro-level demand allocation, which can be biased when electricity use decouples from these indicators (e.g., energy-intensive industrial loads or data centers). Future work will integrate energy-intensity proxies and uncertainty-aware weighting to improve robustness under proxy mismatch and mitigate systematic allocation errors.

CRediT authorship contribution statement

Xuanhao Mu: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Jakob Geiges:** Writing – review & editing, Methodology; **Nan Liu:** Writing – review & editing, Methodology; **Thorsten Schlachter:** Writing – review & editing, Supervision; **Veit Hagenmeyer:** Writing – review & editing, Supervision, Resources, Funding acquisition.

Declaration of competing interest

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

Data availability

Data will be made available on request.

References

- [1] X. Mu, J. Geiges, J. Liu, T. Schlachter, V. Hagenmeyer, Improving spatial allocation for energy system coupling with clustering-based Voronoi diagrams, in: 2025 IEEE 13th International Conference on Smart Energy Grid Engineering (SEGE), 2025. in press.
- [2] K.-K. Cao, et al., Bridging granularity gaps to decarbonize large-scale energy systems—the case of power system planning, *Energy Sci. Eng.* 9 (8) (2021) 1052–1060. <https://doi.org/10.1002/ese3.891>
- [3] M.G. Prina, et al., Classification and challenges of bottom-up energy system models - a review, *Renew. Sustain. Energy Rev.* 129 (2020) 109917. <https://doi.org/10.1016/j.rser.2020.109917>
- [4] M.M. Frysztacki, J. Hörsch, V. Hagenmeyer, T. Brown, The strong effect of network resolution on electricity system models with high shares of wind and solar, *Appl. Energy* 291 (2021) 116726. <https://doi.org/10.1016/j.apenergy.2021.116726>
- [5] R. Martínez-Gordón, G. Morales-España, J. Sijm, A. Faaij, A review of the role of spatial resolution in energy systems modelling: lessons learned and applicability to the north sea region, *Renew. Sustain. Energy Rev.* 141 (2021) 110857. <https://doi.org/10.1016/j.rser.2021.110857>
- [6] Y. Qiu, X. Zhao, D. Fan, S. Li, Y. Zhao, Disaggregating population data for assessing progress of SDGs: methods and applications, *Int. J. Digit. Earth* 15 (2022) 2–29. <https://doi.org/10.1080/17538947.2021.2013553>
- [7] J.K. Wright, A method of mapping densities of population: with cape cod as an example, *Geogr. Rev.* 26 (1) (1936) 103–110.
- [8] S. Patil, N. Pflugradt, J. Weinand, D. Stolten, J. Kropp, A systematic review of spatial disaggregation methods for climate action planning, *Energy AI* (2024). <https://doi.org/10.1016/j.egyai.2024.100386>
- [9] W.C.M. van Beers, J.P.C. Kleijnen, Kriging interpolation in simulation: a survey, in: *Proceedings of the 2004 Winter Simulation Conference, 2004.*, vol. 1, 2004, p. 121. <https://doi.org/10.1109/WSC.2004.1371308>
- [10] L. Hülk, L. Wienholt, I. Cußmann, U.P. Müller, C. Matke, E. Kötter, Allocation of annual electricity consumption and power generation capacities across multiple voltage levels in a high spatial resolution, *Int. J. Sustain. Energy Plan. Manag.* 13 (2017) 79–92. <https://doi.org/10.5278/ijsep.2017.13.6>
- [11] D.Z. Chen, Z. Huang, Y. Liu, J. Xu, On clustering induced Voronoi diagrams, in: 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, 2013, pp. 390–399. <https://doi.org/10.1109/FOCS.2013.49>
- [12] F. Jin, R. Li, H. Wu, Graph neural network-based similarity relationship construction model for geospatial services, *Geo-spat. Inf. Sci.* 27 (2023) 1509–1523. <https://doi.org/10.1080/10095020.2023.2273820>
- [13] K.-H.N. Bui, J. Cho, H. Yi, Spatial-temporal graph neural network for traffic forecasting: an overview and open research issues, *Appl. Intell.* 52 (2021) 2763–2774. <https://doi.org/10.1007/s10489-021-02587-w>
- [14] S. Zhang, H. Tong, J. Xu, R. Maciejewski, Graph convolutional networks: a comprehensive review, *Comput. Soc. Netw.* 6 (2019). <https://doi.org/10.1186/s40649-019-0069-y>
- [15] Z. Hu, Y. Dong, K. Wang, Y. Sun, Heterogeneous graph transformer, *Proc. Web Conf.* (2020). <https://doi.org/10.1145/3366423.3380027>
- [16] Y. Zhou, C. Essayeh, T. Morstyn, Datasets of great Britain primary substations integrated with household heating information, *Data Br.* 54 (2024) 110483. <https://doi.org/10.1016/j.dib.2024.110483>
- [17] O.f.N. Statistics, International territorial level, 2024, Accessed: 2025-08-27, <https://www.ons.gov.uk/methodology/geography/ukgeographies/eurostat>
- [18] K. Jordahl, J.V.d. Bossche, M. Fleischmann, J. Wasserman, J. McBride, J. Gerard, et al., geopandas/geopandas: v0.8.1, 2020, <https://doi.org/10.5281/zenodo.3946761>
- [19] OpenStreetMap contributors, Planet dump retrieved from <https://planet.osm.org>, 2017, <https://www.openstreetmap.org>.
- [20] G. Boeing, Modeling and analyzing urban networks and amenities with OSMnx, 2024, <https://geoffboeing.com/publications/osmnx-paper/>.
- [21] Office for National Statistics, Estimates of the population for england and wales, 2024, Accessed: 2025-04-04, <https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates>.
- [22] Office for National Statistics, Regional gross value added (Balanced) by industry: all ITL regions, 2024, Accessed: 2025-04-04, <https://www.ons.gov.uk/economy/grossvalueaddedgva>.