

Gated Graph Attention Networks for Predicting Duration of Large Scale Power Outages Induced by Natural Disasters

Chenghao Duan*, Chuanyi Ji*, Anwar Walid[†], Scott Ganz[‡]

*School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA
Email: cduan8@gatech.edu, jichuanyi@gatech.edu

[†]Department of Electrical Engineering, Columbia University, New York, NY, USA
Email: aie13@columbia.edu

[‡]School of Business, University of California, Riverside, Riverside, CA, USA
Email: scott.ganz@ucr.edu

Abstract—The occurrence of large-scale power outages induced by natural disasters has been on the rise in a changing climate. Such power outages often last extended durations, causing substantial financial losses and socioeconomic impacts to customers. Accurate estimation of outage duration is thus critical for enhancing the resilience of energy infrastructure under severe weather. We formulate such a task as a machine learning (ML) problem with focus on unique real-world challenges: high-order spatial dependency in the data, a moderate number of large-scale outage events, heterogeneous types of such events, and different impacts in a region within each event. To address these challenges, we develop a Bimodal Gated Graph Attention Network (BiGGAT), a graph-based neural network model, that integrates a Graph Attention Network (GAT) with a Gated Recurrent Unit (GRU) to capture the complex spatial characteristics. We evaluate the approach in a setting of inductive learning, using large-scale power outage data from six major hurricanes in the Southeastern United States. Experimental results demonstrate that BiGGAT achieves a superior performance compared to benchmark models.

Index Terms—Extreme Climate, Graph Neural Network, Gated Recurrent Unit, Machine learning, Power outage, Resilience

I. INTRODUCTION

Natural disasters such as hurricanes, winter storms, wildfires, and flooding have become more frequent and intense in a changing climate [1], [2]. Energy infrastructure is particularly vulnerable, as severe weather events have induced widespread power outages in the US and worldwide [3], [4]. Enhancing the resilience for the energy infrastructure has been advocated by both research communities and industry, where a pertinent aspect of resilience is to reduce interruption duration from power outages [5]. In this context, accurately predicting outage duration is important for emergency preparation and response, resource allocation, and infrastructure planning [6], [7].

We view outage duration prediction as a machine learning (ML) problem: given large-scale power outages induced by prior natural disasters, learn a mapping between outage durations and feature variables on weather, outages, geospatial and demographic characteristics. Importantly, we aim to evaluate how accurately the mapping can predict power

outage durations from an impending weather event. This task presents three real-world challenges: (a) Spatial dependency: Severe weather induced outage durations and feature variables at one location are correlated with those in neighboring and possibly distant areas [8]. (b) Spatial heterogeneity: Weather severity and outages vary spatially, with durations ranging from days near the storm track to hours at locations farther away [7], [9]. (c) Moderate data: Historical data on severe weather-induced outage events is either rare or not publicly available. The data scarcity makes it challenging to train a model that generalizes to new events in different locations.

The objective of this work is to study these challenges to (1) Develop a graph neural network as the mapping, through a novel combination of Graph Attention Network (GAT) [10] with Gated Recurrent Unit (GRU) [11], [12], which captures spatial heterogeneity and dependency in large-scale power outages. (2) Evaluate the model's ability to generalize to unseen events in a setting of inductive learning using field data.

A. Related work

Statistical machine learning approaches on power outages: Prior work on predicting severe weather-induced outage durations has focused on statistical and machine learning methods. Tree-based models achieve state-of-the-art performance for outage duration estimation [13]–[15]. Neural networks [16] are also used for the problem. These approaches, however, do not model spatial characteristics explicitly, and thus require large datasets on historical power outages [15]. When the data are limited, the studies are often restricted to small spatial scopes [13]. Spatial statistics have been applied to outage duration prediction [6], [8]. These methods, however, generally assume homogeneous spatial correlation across locations [17], and thus are unable to characterize spatial heterogeneity in weather-induced power outages.

Graph neural networks and the power grid: Graph based models have been applied to power grid analysis but little to weather-induced large-scale outages using real data. In




Fig. 1. Spatial heterogeneity and dependency of power outages from six hurricanes: (a) Florence, (b) Irma, (c) Laura, (d) Michael, (e) Sally, (f) Zeta. Dark lines: wind boundaries (64, 50, and 34 knots). Clusters on regions with different impacts: Red - significantly impacted counties; Yellow - less impacted counties. Clustered by wind and outage duration.

particular, most prior works focus on power flow prediction, load forecasting, and outage/attack detection [18], [19]. The prior works often rely on synthetic datasets. Field data have mainly been used for outage detection rather than duration prediction [20].

Overall, the research challenges described in the introduction are open and unique for learning from real-world natural disaster-induced large-scale power outage using field data.

II. METHOD

We now describe the development of BiGGAT model. We first present variables and datasets, then problem formulation, algorithm design, and experiment setting.

A. Variables

We select the outage duration as the output variable $\mathbf{Y}$ of BiGGAT model. Here the outage duration $\mathbf{Y}$ characterizes the impact of severe weather-induced large-scale power failures. The feature variables $\mathbf{X}$ include (a) hurricane wind swath as an exogenous cause of power outages, (b) the maximum number of affected customers at given locations, (c) geo-spatial and demographic characteristics of affected regions including spatial adjacency, population density and area, and (d) socioeconomic vulnerability of affected communities.

B. Field data

The power outage data we obtain are from the top six major Atlantic hurricanes that made landfall between 2017 and 2021 in the Southeastern United States (Fig. 1) [6]. These

events caused outages across 557 counties and affected 15 million customers. County-level, time-varying outage data from PowerOutage.com [21] provides the number of customers without service. Outage duration of an affected county is defined as the time from the maximum number of disrupted customers to when this number falls below 5% of customers served.

Weather variables are sourced from the National Hurricane Center's best-track data [22], specifying hurricane wind swaths at each affected location. Maximum sustained wind speeds are sorted into three categories (34-49 knots, 50-63 knots, and 64+ knots). Population, area, and spatial adjacency are obtained from the US Census Bureau [23]. Socioeconomic vulnerability is characterized by the CDC/ATSDR Social Vulnerability Index (SVI) [24], which is widely used in prior works [6].

In practice, outage impact is evaluated by duration levels rather than precise length, e.g., whether outages are short (lasted for hours) or prolonged (for days) [25]. Thus, we label outage durations to three classes, following Department of Energy (DOE)'s restoration guidelines [26]: short (< 2 days), medium (2-6 days), and long (> 6 days). Outage labels correspond to our model output $\mathbf{Y}$. The labeled dataset is summarized in Table I.

TABLE I
DATA SUMMARY: NUMBER OF COUNTIES FOR THREE CLASSES OF OUTAGE DURATIONS

EVENTS	SHORT	MEDIUM	LONG
FLORENCE	68	21	9
IRMA	158	127	29
LAURA	76	18	17
MICHAEL	127	31	18
SALLY	62	7	1
ZETA	179	44	12

C. Problem formulation

We formulate the problem as node-level classification using the graph neural network, that is, learning a mapping from input features $\mathbf{X}$ to outage-duration labels $\mathbf{Y}$ via a graph structure. We make the following assumptions for our formulation:

- 1) Spatial dependency is considered up to n^{th} order geo-spatial neighbors with $n = 1$ for the nearest and $n > 1$ for farther neighbors.
- 2) Spatial dependency is modeled within each event. Cross-event dependencies are not considered in this work.
- 3) County-level variables are used due to data availability.

We let $G = \{V, E, \mathbf{A}\}$ represent the graph structure of the outage data, where V is a set of affected counties (nodes). E consists of neighboring county pairs (edges). $\mathbf{A}$ is a symmetric adjacency matrix, with $\mathbf{A}_{ij} = 1$ if nodes i and j are neighbors, 0 otherwise. For each node $v \in V$ in the data, $\mathbf{x}_v \in \mathbb{R}^{11}$ represents an input feature vector, including the maximum number of affected customers, county population, county area, four variables for social vulnerability and four for hurricane wind swaths. $\mathbf{X} = \{\mathbf{x}_v\}$ for all $v \in V$ represents the set of all node features. $\mathbf{y}_v \in \mathbb{R}^3$ denotes the node-level 3-class

duration labels (short, medium and long), as defined in Field Data. The problem is formalized as a graph learning task, i.e., to obtain a GNN model M that predicts class labels of outage duration $\mathbf{Y}$ using node features $\mathbf{X}$ and graph structure G ,

$$\mathbf{Y} = M(\mathbf{X}, G).$$

D. Bimodal gated graph attention network

We develop a GNN with two key components: (a) a Bimodal Embedding and (b) a Gated Graph Attention mechanism. The Bimodal Embedding is established to capture heterogeneity through unsupervised learning. The Gated Graph Attention mechanism provides a novel integration of self-attention [10] with GRU [11], [12] for message propagation.

1) *Bimodal graph node embedding*: For each node $v \in V$, Bimodal Embedding maps the input $\mathbf{x}_v$ to a message, $\mathbf{m}_{v,k}$ for $k = 0$, where k is the number of recurrent iterations for the Gated Graph Attention layer. Through unsupervised learning (K-means algorithm), the data on wind swath and duration labels form two spatial clusters as illustrated by Fig. 1. The distinct clusters characterize the heterogeneous impact of a weather event. The embedding structure is determined accordingly by the number of clusters. For instance, the two clusters for this case result in two distinct sets of learnable weights, $\beta^{(l_v)}$ ($l_v \in \{1, 2\}$). This results in a simple linear embedding, where for a given node $v \in V$, the embedded message $\mathbf{m}_{v,0}$ is

$$\mathbf{m}_{v,0} = \beta^{(l_v)} \mathbf{x}_v,$$

where $\mathbf{m}_{v,0}$ is a three-dimensional vector representing the initial message at iteration $k = 0$. The number of clusters can vary to represent the complexity of the model and spatial heterogeneity.

2) *Gated graph attention mechanism*: Gated Graph Attention mechanism combines a GRU with self-attention. In particular, after the Bimodal Embedding layer, each node's initial message $\mathbf{m}_{v,0}$ is aggregated via the attention layer into an intermediate embedding $\mathbf{h}_{v,1}$ of the same dimension as $\mathbf{m}_{v,0}$. The attention layer characterizes the spatial dependency weights. A GRU cell then propagates messages to higher-order neighbors with memory: at each iteration $k \in [1, k_{\max}]$, the GRU updates $\mathbf{m}_{v,k}$ from the intermediate embedding $\mathbf{h}_{v,k}$ and the previous message $\mathbf{m}_{v,k-1}$. Here $k_{\max}$ is a hyperparameter of design. The basic message propagation step is:

$$\begin{aligned} \mathbf{h}_{v,k} &= \text{Attn}(\mathbf{m}_{v,k-1}, N_v), \\ \mathbf{m}_{v,k} &= \text{GRU}(\mathbf{m}_{v,k-1}, \mathbf{h}_{v,k}), \end{aligned}$$

where Attn refers to the self-attention based graph message aggregation [10] among the neighbors N_v , and GRU refers to the Gated Recurrent Unit [11], [12].

The embedded message $\mathbf{m}_{v,k}$ is fed into the Gated Graph Attention mechanism for multiple iterations to propagate the influence of each node to its higher-order neighbors. After a given number $k_{\max}$ of iterations, the output message will be passed to a linear readout layer to translate into output labels $\mathbf{o}_v$ at each node. Fig. 2 illustrates the overall structure of the

BiGGAT model. The overall pipeline of the BiGGAT can be summarized as in Algorithm 1

Algorithm 1 Overall BiGGAT pipeline

1. Curate graph-structured outage data using geo-spatial statistics.
 2. Use K-means to determine the structure of the Bimodal Embedding layer.
 3. Embed node message $\mathbf{m}_{v,0}$ using the Bimodal Embedding layer.
 4. Use the self-attention mechanism to update the hidden embedding $\mathbf{h}_{v,k}$.
 5. Feed both the previous embedded message $\mathbf{m}_{v,k-1}$ and the current hidden embedding $\mathbf{h}_{v,k}$ into GRU for the current embedded message $\mathbf{m}_{v,k}$.
 6. Repeat step 4 and 5 for $k \in [1, k_{\max}]$.
 7. Feed the embedded message from step 6 to a linear read out layer for prediction $\mathbf{o}_v$. Evaluate the loss and accuracy.
-

E. Experiment setting

We adopt the setting of inductive learning to apply BiGGAT to real-world scenarios. We further specify the experiment designs and describe three metrics for performance evaluation.

1) *Inductive learning*: The setting of inductive learning is desirable for real-world implementation where a model is trained on historical events and then predicts the outage durations of an unseen event. To implement such a setting, we hold out one event as the testing set and use the other five events to train BiGGAT. For test samples, the cluster labels l_v are inferred using input features.

One key design for the model is the graph structure of the data - characterized by the neighborhood N_v of each node. To capture sufficient spatial dependency, we define N_v to include all counties up to the highest spatial order of significant correlation. Meanwhile, for simplicity, we define N_v to include neighbors up to the same spatial order n for all nodes within a given event. The actual value of n is obtained from data: we measure the spatial correlation of the feature variable, peak outages, using the global Moran's I [27]. The Moran's I, a widely-used spatial statistic, measures the weighted covariance of the variables among the immediate geo-spatial neighbors. We extend the measure to farther spatial neighbors by adapting the spatial weight matrix A in the Moran's I calculation with an n-hop neighbor adjacency matrix $A_{\text{spatial}}^{(n)}$:

$$A_{\text{spatial}}^{(n)} = I\{A_{\text{spatial}}^n > 0\} - I\{A_{\text{spatial}}^{n-1} > 0\}$$

Where A_{spatial} denotes the spatial adjacency matrix of affected counties and I denotes an indicator function. Thus, $A_{\text{spatial}}^{(n)}$ is nonzero only at entries where two given locations are exactly n^{th} order neighbor to each other without redundancy. Such an n-hop Moran's I effectively captures the correlation among the n-hop spatial neighbors.

The n-hop Moran's I (Table II) shows significant positive spatial correlation for at least the second-order geo-spatial neighbors. Therefore, we select the highest significant n-hop




Fig. 2. Overall BiGGAT structure: Node features are fed into the Bimodal Embedding, then passed through the Gated Graph Attention mechanism. The aggregated messages are transformed to outputs by a Linear Readout layer.

TABLE II
GLOBAL N-HOP MORAN'S I OF PEAK OUTAGE NUMBER

EVENTS	HIGHEST N	GLOBAL N-HOP MORAN'S I
FLORENCE	3	0.273* (6.806)
IRMA	6	0.070* (4.723)
LAURA	5	0.112* (3.499)
MICHAEL	3	0.053* (1.998)
SALLY	2	0.110* (1.864)
ZETA	4	0.154* (6.049)

* denotes rejection of null hypothesis with $p \leq 0.1$.
Values in brackets are z-scores.

number to define the size of neighborhood N_v . In addition, we choose the max recurrence iteration number $kmax$ to be 2. Thus, with control from GRU, the messages are not only aggregated within neighborhood N_v but also propagated to farther neighbors beyond N_v . This design ensures that our model can fully capture the high-order spatial correlations present in each event.

Our model is trained for node-level multi-class classification. Three metrics are used to benchmark the performances: (1) Classification Accuracy, (2) Macro F1 Score, and (3) Balanced Accuracy. We use the cross-entropy loss and ADAM optimizer for training, which are implemented using Pytorch, PyGeometrics, and Sklearn packages via Google Colab.

III. RESULTS

A. Inductive learning model performance

The BiGGAT and benchmark models, including XGBoost (XGB), Random Forest (RF), GAT, and BiGAT, are trained on five events and tested on the remaining unseen event. These experiments are repeated six times, with each event serving as the test set once. For each model, the corresponding hyperparameters are iteratively tuned to the corresponding optimal performances.

Event-wise comparison and averaged metrics are shown in Fig. 3 and Table III respectively. BiGGAT consistently outperforms the tree-based models (XGB and RF), indicating that graph-based models better capture spatial dependencies in the outage data. Among graph-based models, BiGGAT and BiGAT outperforms GAT in five events, underscoring the value




Fig. 3. Model performance comparison for (a) Florence, (b) Irma, (c) Laura, (d) Michael, (e) Sally, and (f) Zeta. Each subplot shows classification accuracy, Macro F1, and balanced accuracy for XGB, RF, GAT, BiGAT, and BiGGAT.

TABLE III
EVENT AVERAGE PERFORMANCE OF MODELS

MODELS	ACCURACY	BALANCED ACC	MACRO F1
XGB	79.7%	67.8%	0.657
RF	81.5%	69.4%	0.682
GAT	83.7%	68.1%	0.678
BiGAT	85.1%	70%	0.686
BiGGAT	85.9%	78.3%	0.762

of modeling spatial heterogeneity via Bimodal Embedding. Compared with BiGAT, which uses only nearest neighbors and without GRU, BiGGAT achieves better class-balanced performance in all events. These improvements indicate better accuracies for medium- and long-duration classes in the imbalanced dataset. There are fewer medium and long duration samples in the data, which are more challenging to learn.

This also highlights the importance of modeling high-order spatial dependencies. Overall, the BiGGAT shows significant improvements over existing models on learning outage durations in the challenging setting of inductive learning.

B. Generalization of GRU and error analysis

The inductive setting for outage duration prediction requires generalization across hurricanes affecting different regions. To evaluate this, we split test counties from unseen events into two groups: those with geographical overlap with the training set (“test with overlap”) and those without any geographical or structural connection (“absolute disjointed”). The latter group contains 265 test counties and represents the most challenging cases for generalization. Nonetheless, their test accuracy (82.4%) is comparable to the overall test accuracy, and errors from these counties account for only 27.3% of all test errors. These results indicate that BiGGAT generalizes well to unseen and structurally unconnected samples.

IV. CONCLUSION

This work presents a novel application of graph neural networks for predicting duration of power outages induced by natural disasters. We develop BiGGAT that captures spatial heterogeneity and high-order spatial dependencies in the field data. BiGGAT learns from and is tested by six historical Atlantic hurricanes under an inductive learning setting. To achieve robustness with moderate event data, the model adopts a relatively simple architecture with Bimodal Embedding, self-attention and GRU modules for message passing. Experimental results show that BiGGAT significantly outperforms benchmark models.

Further, the Gated Graph Attention mechanism enhances the generalization performance and class-balanced metrics in inductive learning. The model also performs robustly on the test samples that are located in differently regions from the training data. These results underscore the importance of modeling higher-order spatial dependencies, particularly under severe weather events, and validate the use of GRU and self-attention for message passing.

There are several research directions for future study: (a) extending BiGGAT to other events and disaster types, including wildfires, winter storms and tornadoes, (b) adapting the model to data with finer spatial resolution, e.g., private data owned by power utilities, (c) exploring how predicted outage impacts benefit vulnerability assessments and infrastructure. All these will benefit informed decision-making in disaster preparation and resilience strategies.

REFERENCES

- [1] S. Seneviratne, X. Zhang, M. Adnan, W. Badi, C. Dereczynski, A. Di Luca, S. Ghosh, I. Iskandar, J. Kossin, S. Lewis, F. Otto, I. Pinto, M. Satoh, S. Vicente-Serrano, M. Wehner, and B. Zhou, *Weather and Climate Extreme Events in a Changing Climate*. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press, 2021, p. 1513–1766.
- [2] K. Emanuel, “Atlantic tropical cyclones downscaled from climate re-analyses show increasing activity over past 150 years,” *Nature communications*, vol. 12, no. 1, pp. 1–8, 2021.
- [3] A. M. Stanković, K. L. Tomsovic, F. De Caro, M. Braun, J. H. Chow, N. Čukalevski, I. Dobson, J. Eto, B. Fink, C. Hachmann *et al.*, “Methods for analysis and quantification of power system resilience,” *IEEE Transactions on Power Systems*, vol. 38, no. 5, pp. 4774–4787, 2022.
- [4] A. Adhikari, C. Wertz, A. Dubey, A. Ahmad, and I. Dobson, “Quantifying power systems resilience using statistical analysis and bayesian learning,” *arXiv preprint arXiv:2511.03043*, 2025.
- [5] L. Xu, K. Feng, N. Lin, A. Perera, H. V. Poor, L. Xie, C. Ji, X. A. Sun, Q. Guo, and M. O’Malley, “Resilience of renewable power systems under climate risks,” *Nature Reviews Electrical Engineering*, vol. 1, no. 1, pp. 53–66, 2024.
- [6] S. C. Ganz, C. Duan, and C. Ji, “Socioeconomic vulnerability and differential impact of severe weather-induced power outages,” *PNAS nexus*, vol. 2, no. 10, p. pgad295, 2023.
- [7] A. H. Afsharinejad, C. Ji, and R. Wilcox, “Large-scale data analytics for resilient recovery services from power failures,” *Joule*, vol. 5, no. 9, pp. 2504–2520, 2021.
- [8] K. Best, S. Kerr, A. Reilly, A. Patwardhan, D. Niemeier, and S. Guikema, “Spatial regression identifies socioeconomic inequality in multi-stage power outage recovery after hurricane isaac,” *Natural Hazards*, vol. 117, no. 1, pp. 851–873, 2023.
- [9] C. Ji, Y. Wei, H. Mei, J. Calzada, M. Carey, S. Church, T. Hayes, B. Nugent, G. Stella, M. Wallace *et al.*, “Large-scale data analysis of power grid resilience across multiple us service regions,” *Nature Energy*, vol. 1, no. 5, pp. 1–8, 2016.
- [10] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio, “Graph attention networks,” *arXiv preprint arXiv:1710.10903*, 2017.
- [11] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel, “Gated graph sequence neural networks,” *arXiv preprint arXiv:1511.05493*, 2015.
- [12] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, “Learning phrase representations using rnn encoder-decoder for statistical machine translation,” *arXiv preprint arXiv:1406.1078*, 2014.
- [13] R. Nateghi, S. D. Guikema, and S. M. Quiring, “Forecasting hurricane-induced power outage durations,” *Natural hazards*, vol. 74, pp. 1795–1811, 2014.
- [14] D. B. McRoberts, S. M. Quiring, and S. D. Guikema, “Improving hurricane power outage prediction models through the inclusion of local environmental factors,” *Risk analysis*, vol. 38, no. 12, pp. 2722–2737, 2018.
- [15] J. Fang, H. Wang, F. Yang, K. Yin, X. Lin, and M. Zhang, “A failure prediction method of power distribution network based on pso and xgboost,” *Australian journal of electrical and electronics engineering*, vol. 19, no. 4, pp. 371–378, 2022.
- [16] J. Yu, Z. Wang, A. Majumdar, and R. Rajagopal, “Deepsolar: A machine learning framework to efficiently construct a solar deployment database in the united states,” *Joule*, vol. 2, no. 12, pp. 2605–2617, 2018.
- [17] J. LeSage and R. K. Pace, *Introduction to spatial econometrics*. Chapman and Hall/CRC, 2009.
- [18] W. Liao, B. Bak-Jensen, J. R. Pillai, Y. Wang, and Y. Wang, “A review of graph neural networks and their applications in power systems,” *Journal of Modern Power Systems and Clean Energy*, vol. 10, no. 2, pp. 345–360, 2021.
- [19] A. K. Dey, Y. R. Gel, and H. V. Poor, “What network motifs tell us about resilience and reliability of complex networks,” *Proceedings of the National Academy of Sciences*, vol. 116, no. 39, pp. 19368–19373, 2019.
- [20] D. Owerko, F. Gama, and A. Ribeiro, “Predicting power outages using graph neural networks,” in *2018 IEEE Global Conference on Signal and Information Processing (globalsip)*. IEEE, 2018, pp. 743–747.
- [21] PowerOutages.US, “Power outage records of the Southeast US,” 2021.
- [22] National Hurricane Center, “The revised atlantic hurricane database (hurdat2),” 2023.
- [23] Census Bureau, “American community survey 1-year estimates,” 2020.
- [24] B. E. Flanagan, E. J. Hallisey, E. Adams, and A. Lavery, “Measuring community vulnerability to natural and anthropogenic hazards: the centers for disease control and prevention’s social vulnerability index,” *Journal of environmental health*, vol. 80, no. 10, p. 34, 2018.
- [25] Georgia Power, *Georgia Power Plants Emergency Action Plan*, 2024.
- [26] Department of Energy, *Energy Emergency Response Playbook for States and Territories*, May 2022.
- [27] P. A. Moran, “Notes on continuous stochastic phenomena,” *Biometrika*, vol. 37, no. 1/2, pp. 17–23, 1950.