



LAWRENCE  
LIVERMORE  
NATIONAL  
LABORATORY

LLNL-CONF-794857

# TOWARDS AN EFFICIENT AND GENERAL FRAMEWORK OF ROBUST TRAINING FOR GRAPH NEURAL NETWORKS

K. Xu, S. Liu, P. Chen, M. Sun, C. Ding, B. Kailkhura, X. Lin

October 22, 2019

International Conference on Acoustics, Speech, and Signal  
Processing  
Barcelona, Spain  
May 4, 2020 through May 8, 2020

## **Disclaimer**

---

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

# TOWARDS AN EFFICIENT AND GENERAL FRAMEWORK OF ROBUST TRAINING FOR GRAPH NEURAL NETWORKS

*Kaidi Xu<sup>1</sup>, Sijia Liu<sup>2</sup>, Pin-Yu Chen<sup>2</sup>, Mengshu Sun<sup>1</sup>, Caiwen Ding<sup>3</sup>, Bhavya Kailkhura<sup>4</sup>, Xue Lin<sup>1</sup>*

<sup>1</sup>Electrical & Computer Engineering, Northeastern University, Boston, USA

<sup>2</sup>MIT-IBM Watson AI Lab, IBM Research

<sup>3</sup>Computer Science and Engineering, University of Connecticut, Storrs, USA

<sup>4</sup>Lawrence Livermore National Laboratory, Livermore, USA

## ABSTRACT

Graph Neural Networks (GNNs) have made significant advances on several fundamental inference tasks. As a result, there is a surge of interest in using these models for making potentially important decisions in high-regret applications. However, despite GNNs' impressive performance, it has been observed that carefully crafted perturbations on graph structures (or nodes attributes) lead them to make wrong predictions. Presence of these adversarial examples raises serious security concerns. Most of the existing robust GNN design/training methods are only applicable to white-box settings where model parameters are known and gradient based methods can be used by performing convex relaxation of the discrete graph domain. More importantly, these methods are not efficient and scalable which make them infeasible in time sensitive tasks and massive graph datasets. To overcome these limitations, we propose a general framework which leverages the greedy search algorithms and zeroth-order methods to obtain robust GNNs in a generic and an efficient manner. On several applications, we show that the proposed techniques are significantly less computationally expensive and, in some cases, more robust than the state-of-the-art methods making them suitable to large-scale problems which were out of the reach of traditional robust training methods.

**Index Terms**— Graph neural networks, adversarial training, robustness, greedy algorithm, large-scale learning

## 1. INTRODUCTION

Graph structured representations are one of the most commonly encountered data structure that naturally arises in nearly all scientific and engineering application [1]. With the widespread use of networks neuroscience, molecular chemistry and other fields, it is not surprising that machine learning on graph data has become a key learning tool. Specifically, Graph neural networks (GNNs) have made significant advances on several fundamental tasks ranging from node classification to graph classification [2, 3]. Despite GNNs' impressive performance on inferring from graph data, their

susceptibility to test-time adversarial examples (i.e., carefully crafted perturbations to fool these models) is a major hurdle in a universal acceptance of GNN solutions in several high-regret applications. These results are consistent with the adversarial attacks on audios and images [4, 5, 6, 7, 8, 9].

Recently, a few attempts have been carried out in the robust machine learning community to devise robust GNN training methods, e.g., adversarial training. Existing robust training methods for graph data [10, 11] utilize convex relaxation to make the search domain continuous so that first-order methods can be used. Unfortunately, these approaches are highly inefficient (or time consuming) due to additional computations involved in the relaxation procedure and are infeasible to large-scale graphs. Furthermore, in practice the gradient information may not even be available, e.g., either due to inaccessible of the model parameters or model itself is discrete and non-continuous. In such scenarios, existing robust training methods fall short. To overcome these drawbacks, this paper proposes a general framework which leverages the greedy search algorithms and zeroth-order methods to obtain robust GNNs. Specifically, to address the inefficiency/scalability issue, we propose Greedy Topology Attacking method so that gradient-based adversarial training becomes plausible for massive size graphs. Next, to address the unavailability of gradient information issue, we propose Zeroth-order based Greedy Attack for gradient-free adversarial training. Benchmarking on node classification tasks using GNNs, our Greedy Topology Attack method can achieve similar performance with current state-of-the-art attacks with significant speed up subject to the same topology perturbation budget. This demonstrates the effectiveness of our attack generation method through the lens of greedy search algorithm. Next, by leveraging our proposed greedy topology attack, we benchmark the robustness of our adversarial training technique for GNNs under different attacks and applications. We show that the proposed adversarial training technique is significantly less computationally expensive and, in some cases, more robust than the state-of-the-art methods making them suitable to large-scale problems.

## 2. RELATED WORKS

Some recent attentions have been paid to the vulnerability of GNNs. The authors in [12] investigated test-time non-targeted adversarial attacks on both graph classification and node classification problem by adding or deleting edges from a graph. The authors in [13] considered both test-time (evasion) and training-time (data poisoning) attacks on node classification task. Different with [12], the node attributes were allowed to modify in [13]. Their algorithm is for targeted attacks on single node. It was shown that small perturbations on the graph structure and node features are able to misclassified the target node. In [14], training-time attacks on GNNs were also studied for node classification by perturbing the graph structure. The authors solved a min-max problem in training-time attacks using meta-gradients and treated the graph topology as a hyper-parameter to optimize.

On the defense side, very recent work [11] proposed adversarial training on GNNs, however, this method needs to perform convex relaxation, probabilistic sampling and bisection which pay extra computational cost and is not feasible to large graph datasets. Different from the adversarial training, [10] considered the problem of certifiable robustness and also achieved robust model under a pre-specific bound.

## 3. PROBLEM STATEMENT

A GNN predicts the class of an unlabeled node under the graph topology. The  $k$ -th layer computation in a GNN can be formalated as

$$H^{(k)} = \sigma \left( \tilde{A} H^{(k-1)} (W^{(k-1)}) \right), \quad (1)$$

where  $H^{(1)} = X$  is the node feature matrix. Here  $\tilde{A}$  is defined as a normalized adjacency matrix  $\tilde{A} = \hat{D}^{-1/2} \hat{A} \hat{D}^{-1/2}$ , where  $\hat{A} = A + I$  and  $A \in \{0, 1\}^{N \times N}$  refers to the adjacency matrix. We refer readers to [2] for more details about GNN.  $\sigma$  is the ReLU activation function and  $W$  is weight matrices.

### 3.1. Adversarial Training for Graphs

Following [15], we consider an adversarial variant of standard Empirical Risk Minimization (ERM), where we minimize the risk over adversarial examples:

$$\underset{W}{\text{minimize}} \underset{A' \in \mathcal{C}}{\text{maximize}} \mathcal{L}(A', W; X, y_L). \quad (2)$$

The training loss  $\mathcal{L}$  is cross-entropy error over all training data  $(X, y_L)$ .  $A'$  is the perturbed adjacency matrix by the adversarial attack. Here  $\mathcal{C}$  are the constrains such as the maximum number of edges that can be perturbed,  $A'$  should be symmetric and most importantly,  $A' \in \{0, 1\}^{N \times N}$ .

Intuitively, adversarial training injects adversarial examples into training data via inner maximization problem to increase robustness of outer risk minimization (or training

problem. The performance of adversarial training heavily depends on (a) the quality of the adversarial examples, and (b) the efficiency/scalability of the attack algorithms. Note that discrete constraint  $\mathcal{C}$  in problem (3) make the gradient descent inapplicable. As mentioned before, current approaches to handle these constraints are computationally inefficient in solving inner optimization problem. Therefore, we first propose efficient algorithms to generate adversarial attacks via greedy algorithms.

### 3.2. Efficient Adversarial Example Generation

In this section, we propose two approaches to generate adversarial examples based on greedy algorithms and zeroth-order methods.

#### 3.2.1. Greedy Topology Attack

We fix node attributes  $X$  and only consider edge perturbations as the adversarial attack. We find the perturbed adjacency matrix  $A'$  to minimize the negative of training loss

$$\begin{aligned} & \underset{A'}{\text{minimize}} \quad -\mathcal{L}(f(A'; W, X, y_L)) \\ & \text{subject to} \quad A' \in \mathcal{C}. \end{aligned} \quad (3)$$

Next, we introduce our Greedy Topology Attack (GTA) in Algorithm 1 which is able to handle discrete constraints.

---

#### Algorithm 1 Greedy topology attack (GTA)

---

```

1: Input: graph  $G = (A, X)$ , number of maximum edges
   can be changed  $M$ , greedy search step  $n$ , label  $y_L$ 
2: Output: Modified adversarial adjacency matrix  $A'$ 
3:  $A' \leftarrow A$ 
4:  $R \leftarrow \emptyset$ 
5: while  $\|A' - A\|_0 < 2M$  do
6:    $\triangleright$  Compute gradient of Eq. (3)
7:    $g_{A'} \leftarrow \nabla_{A'} \mathcal{L}(f(A'; W, X, y_L))$ 
8:    $S \leftarrow \nabla_{A'} \odot (-2A' + 1)$ 
9:    $e \leftarrow$  top  $n$  elements in  $S$  if they are not in  $R$ 
10:   $A' \leftarrow$  remove or insert edge  $e$  from/to  $A'$ 
11:   $R \leftarrow$  add  $e$  in record  $R$ 
12: end while
13: return:  $A'$ 

```

---

To satisfy the symmetric attribute of  $A$ , we only perturb the upper triangular part of the matrix  $A$  and replicate it to the lower triangular part. This is the reason that in Line 4 we have  $2M$ . In Line 6, we compute the gradient over  $A'$  of Eq. (3). In Line 7, we conduct element-wise product between the gradient  $g_{A'}$  and a flipped adjacency matrix (0 becomes 1, and 1 becomes -1). This step actually help us express the edge that can be removed or inserted from/to  $A'$  also with the gradient value. In Line 8, we speed up the greedy search process. By calculating the  $g_{A'}$  once, we select at most  $n$  edges to be changed in  $A'$ , this improvement accelerate the whole

attacking method around  $n$  times. The sensitivity between the attacking performance and choice of  $n$  will also be shown in experiment part. Line 9 records all the modified edges to avoid sometimes change repeatedly.

Overall, our attack method utilize the gradient information but also satisfies the strict constraints compared to the conventional gradient descent methods. The proposed greedy search step also helps in terms of the convergence speed.

### 3.2.2. Zeroth-Order Greedy Topology Attack

Note although during training stage, we usually use sparse matrix to store  $A$ , when optimizing graph over  $A'$ ,  $g_{A'}$  may be very dense, so that even efficient method like GTA may not be feasible to extremely large graph. Thus, to further improve the efficiency of GTA, we also address the expensiveness of obtaining gradient issue via zeroth-order methods. As an extension to our GTA method, we propose Zeroth-order Greedy Topology Attack (ZO-GTA) to improve the feasibility on large graph. We summarize our ZO-GTA in Algorithm 2.

#### Algorithm 2 Greedy topology attack (GTA)

```

1: Input: graph  $G = (A, X)$ , number of maximum edges
   can be changed  $M$ , greedy search step  $n$ , label  $y_L$ 
2: Output: Modified adversarial adjacency matrix  $A'$ 
3:  $A' \leftarrow A$ 
4:  $L^{(0)} \leftarrow$  training loss  $\mathcal{L}(A'; W, X, y_L)$ 
5:  $t \leftarrow 1$ 
6: while  $\|A' - A\|_0 < 2M$  do
7:    $A' \leftarrow$  random choice  $n$  elements in  $A'$  and flip them.
8:    $L^t \leftarrow$  training loss  $\mathcal{L}(A'; W, X, y_L)$ 
9:   if  $L^t > L^{t-1}$  then
10:    continue
11:   else
12:     $A' \leftarrow$  flip again the  $n$  elements in  $A'$ 
13:   end if
14: end while
15: return:  $A'$ 

```

Intuitively, we continuously check randomly flipped  $n$  nodes to see whether they can help us maximize  $\mathcal{L}$  or not. If yes, we keep the  $n$  nodes flipped, otherwise, we try this procedure again. The ZO-GTA evade to calculate the gradients in problem 3 and preserve the discrete nature of  $A'$  as well.

### 3.3. Algorithm for GNN Adversarial Training

With the aid of our effective attack methods, the robust training for GNNs (see (2)) will be solved in this section.

Note that the inner maximization problem in (2) is exactly same as Eq. (3). Thus, our GTA methods can approximate this inner optimization. For the outer minimize function, we follow [2] by using gradient descent to update  $W$ . This formulation aims to minimize the training loss for the

worst case topology perturbations. We summarize our robust training method for GNNs in Algorithm 3.

---

#### Algorithm 3 Robust training for solving problem (2)

---

```

1: Input: graph  $G = (A, X)$ , number of maximum edges
   can be changed  $N$ , greedy search step  $n$ , learning rates
    $\beta$ , and iteration numbers  $T$ , label  $y_L$ 
2: Output:  $W$ 
3: randomly initialize  $W^{(0)}$ 
4: for  $t = 1, 2, \dots, T$  do
5:   inner minimization over  $A'$ : given  $W^{(t-1)}$ , running
      GTA (3) and obtain  $A'$ 
6:   outer maximization over  $W$ : given  $A'$ , obtain
       $W^t = W^{t-1} - \beta \nabla_W \mathcal{L}(A', W^{t-1}; X, y_L)$ 
7: end for
8: return  $W$ 

```

---

Note that this robust training algorithm is general enough to accommodate any existing attack method as long as constraint  $A' \in \mathcal{C}$  in the inner minimization step is satisfied. For example, one can replace GTA with ZO-GTA or CE-PGD attack [11] in Algorithm 3 to get different variants of robust models. Different attack methods will results in robust models with different degrees of robustness.

## 4. EXPERIMENTS

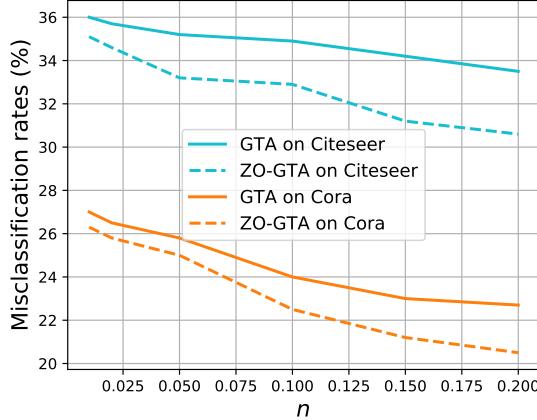
### 4.1. Experimental Setup

In this section, our experimental results are presented for both GTA and robust training on graph convolutional networks (GCN) [2] on three popular datasets: Cora, Citeseer and Pubmed [16]. The datasets statistics are summarized in Table 1.

**Table 1.** Dataset statistics summary [17]

Dataset	Nodes ( $N$ )	Edges	Classes	Features
Cora	2,708	5,429	7	1,433
Citeseer	3,327	4,732	6	3,703
Pubmed	19,717	44,338	3	500

We demonstrate the misclassification rate (namely, 1-prediction accuracy on unseen test nodes) of the proposed GTA and ZO-GTA method. To provide reliable results, we repeat these experiments 5 times based on different splits of training/testing nodes and report  $mean \pm standard\ deviation$  of the performance. We follow [14, 11] where for test nodes' predicted labels (not their ground-truth label, generate by an independent pre-trained model) can be used during the attack.



**Fig. 1.** Misclassification rate (in %) of GTA/ZO-GTA on Citeseer/Cora datasets vs  $n$ .  $n$  is choosing from  $0.01N$  to  $0.2N$ .

## 4.2. Attack Performance

As we mentioned in Section 3.2.1 and 3.2.2, the hyper parameter  $n$  affects the speed and the performance of our final model. First we fix  $M = 0.05N$  in each dataset. Next, we conduct the experiment by searching  $n \in [0.01N, 0.2N]$  which means we attack a graph by around 5 to 100 iterations. The results are shown in Figure 1. As we can see, the misclassification rate drop as  $n$  increases. However, the smaller the  $n$  is, more iterations we need to execute the attack. Also, we can conclude that GTA preforms better than ZO-GTA. This is intuitive as ZO-GTA is an approximate scheme as compared to GTA which use gradient information. Therefore, in the following experiments we set  $n = 0.05N$ . We compare our GTA and ZO-GTA methods with DICE ('delete edges internally, connect externally') [18], CE-PGD and CW-PGD [11]. We follow the hyper parameters and experimental settings as given in in [11] for a fair comparison. The attack results and run time are reported in Table 2 and Table 3, respectively.

**Table 2.** Misclassification rates (in %) of GTA, ZO-GTA, DICE, CE-PGD and CW-PGD over Cora Citeseer and Pubmed datasets. No result means that the method is not feasible for this dataset.)

	Cora	Citeseer	Pubmed
Clean	$18.2 \pm 0.1$	$28.9 \pm 0.3$	$16.9 \pm 0.5$
GTA	$25.8 \pm 0.1$	$35.2 \pm 0.1$	<b><math>19.3 \pm 0.4</math></b>
ZO-GTA	$24.9 \pm 0.3$	$33.1 \pm 0.3$	$18.8 \pm 0.5$
DICE	$18.9 \pm 0.3$	$29.8 \pm 0.4$	$17.4 \pm 0.6$
CE-PGD	<b><math>28.0 \pm 0.1</math></b>	$36.0 \pm 0.2$	-
CW-PGD	$27.8 \pm 0.4$	<b><math>37.1 \pm 0.5</math></b>	-

It can be clearly seen that GTA and ZO-GTA perform competitively with gradient methods, i.e., CE-PGD and CW-

PGD and yield significantly better computational efficiency.

**Table 3.** Running time (in seconds) of GTA, ZO-GTA, DICE, CE-PGD and CW-PGD over Cora Citeseer and Pubmed datasets. No result means that the method is not feasible for this dataset.

	Cora	Citeseer	Pubmed
GTA	$37 \pm 1$	$36 \pm 1$	$109 \pm 5$
ZO-GTA	<b><math>19 \pm 3</math></b>	<b><math>21 \pm 3</math></b>	$43 \pm 5$
CE-PGD	$147 \pm 4$	$144 \pm 4$	-
CW-PGD	$151 \pm 5$	$142 \pm 2$	-

## 4.3. Defense Performance

We next show the improved robustness of GCN by leveraging our proposed robust training algorithm against different topology attacks. We set  $T = 1000$ ,  $\beta = 0.01$  and  $M = 0.05N$ . We evaluate Algorithm 3 on Cora dataset to compare it with [11]. GTA is used to solve the inner maximization problem since it yields better misclassification rate compared to ZO-GTA. To test the robustness of the models, we use both GTA and CE-PGD attacking methods to attack them. Results in Table 4 show that the GTA based robust training is competitive with CE-PGD based robust training. However, recall the results in Table 3, GTA is much more computationally efficient than CE-PGD which helps our GTA method to handle time sensitive tasks and massive graphs.

**Table 4.** Misclassification rates (in %) of two robust training methods against GTA and CE-PGD attack on Cora dataset. 'clean-' means test without attack.

	GTA training	CE-PGD training
clean-nature		$18.2 \pm 0.1$
clean-robust	$18.1 \pm 0.2$	$18.0 \pm 0.3$
GTA attack	$20.4 \pm 0.3$	$20.8 \pm 0.4$
CE-PGD attack	$22.7 \pm 3$	$22.0 \pm 0.2$

## 5. CONCLUSION

In this paper, we first introduce two GNN attacking method GTA and ZO-GTA based on the greedy search and zeroth-order algorithms. The proposed approaches are shown to be competitive with state-of-the-art attacks. Next a general and efficient framework of robust training on graph neural networks is proposed. Our experimental results show that the proposed robust training method is significantly computationally less expensive while achieving high robustness to various adversarial attacks. This makes them an potentially viable candidate to handle large-scale and time sensitive problems.

## 6. REFERENCES

[1] Fernando Gama, Antonio G Marques, Geert Leus, and Alejandro Ribeiro, “Convolutional neural network architectures for signals supported on graphs,” *IEEE Transactions on Signal Processing*, vol. 67, no. 4, pp. 1034–1049, 2018.

[2] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” *arXiv preprint arXiv:1609.02907*, 2016.

[3] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka, “How powerful are graph neural networks?,” in *International Conference on Learning Representations*, 2019.

[4] Sijia Liu, Bhavya Kailkhura, Pin-Yu Chen, Paishun Ting, Shiyu Chang, and Lisa Amini, “Zeroth-order stochastic variance reduction for nonconvex optimization,” in *Advances in Neural Information Processing Systems*, 2018, pp. 3727–3737.

[5] Thomas A Hogan and Bhavya Kailkhura, “Universal hard-label black-box perturbations: Breaking security-through-obscenity defenses,” *arXiv preprint arXiv:1811.03733*, 2018.

[6] Pu Zhao, Sijia Liu, Pin-Yu Chen, Nghia Hoang, Kaidi Xu, Bhavya Kailkhura, and Xue Lin, “On the design of black-box adversarial examples by leveraging gradient-free optimization and operator splitting method,” *arXiv preprint arXiv:1907.11684*, 2019.

[7] Kaidi Xu, Sijia Liu, Gaoyuan Zhang, Mengshu Sun, Pu Zhao, Quanfu Fan, Chuang Gan, and Xue Lin, “Interpreting adversarial examples by activation promotion and suppression,” *arXiv preprint arXiv:1904.02057*, 2019.

[8] Kaidi Xu, Sijia Liu, Pu Zhao, Pin-Yu Chen, Huan Zhang, Quanfu Fan, Deniz Erdogmus, Yanzhi Wang, and Xue Lin, “Structured adversarial attack: Towards general implementation and better interpretability,” in *International Conference on Learning Representations*, 2019.

[9] Hongge Chen, Huan Zhang, Pin-Yu Chen, Jinfeng Yi, and Cho-Jui Hsieh, “Attacking visual language grounding with adversarial examples: A case study on neural image captioning,” in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2018, pp. 2587–2597.

[10] Daniel Zügner and Stephan Günnemann, “Certifiable robustness and robust training for graph convolutional networks,” in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 2019, pp. 246–256.

[11] Kaidi Xu, Hongge Chen, Sijia Liu, Pin-Yu Chen, Tsui-Wei Weng, Mingyi Hong, and Xue Lin, “Topology attack and defense for graph neural networks: An optimization perspective,” in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2019.

[12] H. Dai, H. Li, T. Tian, X. Huang, L. Wang, J. Zhu, and L. Song, “Adversarial attack on graph structured data,” *arXiv preprint arXiv:1806.02371*, 2018.

[13] D. Zügner, A. Akbarnejad, and S. Günnemann, “Adversarial attacks on neural networks for graph data,” in *KDD*. ACM, 2018.

[14] Daniel Zügner and Stephan Günnemann, “Adversarial attacks on graph neural networks via meta learning,” in *ICLR*, 2019.

[15] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu, “Towards deep learning models resistant to adversarial attacks,” in *International Conference on Learning Representations*, 2018.

[16] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad, “Collective classification in network data,” *AI magazine*, vol. 29, no. 3, pp. 93–93, 2008.

[17] Zhilin Yang, William W Cohen, and Ruslan Salakhutdinov, “Revisiting semi-supervised learning with graph embeddings,” *arXiv preprint arXiv:1603.08861*, 2016.

[18] Marcin Waniek, Tomasz P Michalak, Michael J Wooldridge, and Talal Rahwan, “Hiding individuals and communities in a social network,” *Nature Human Behaviour*, vol. 2, no. 2, pp. 139, 2018.