JOINT LEARNING OF OBJECT GRAPH AND RELATION GRAPH FOR VISUAL QUESTION ANSWERING

Hao Li^{1,†}, Xu Li², Belhal Karimi², Jie Chen^{1,3}, Mingming Sun^{2,*}

¹School of Electronic and Computer Engineering, Peking University, China; ²Cognitive Computing Lab, Baidu Research, China; ³ Peng Cheng Laboratory, Shenzhen, China; lihao1984@pku.edu.cn, chenj@pcl.ac.cn, {lixu13, belhalkarimi, sunmingming01}@baidu.com

ABSTRACT

Modeling visual question answering (VQA) through scene graphs can significantly improve the reasoning accuracy and interpretability. However, existing models answer poorly for complex reasoning questions with attributes or relations, which causes false attribute selection or missing relation in Figure 1(a). It is because these models cannot balance all kinds of information in scene graphs, neglecting relation and attribute information. In this paper, we introduce a novel Dual Message-passing enhanced Graph Neural Network (DM-GNN), which can obtain a balanced representation by properly encoding multi-scale scene graph information. Specifically, we (i) transform the scene graph into two graphs with diversified focuses on objects and relations: Then we design a *dual structure* to encode them, which increases the weights from relations (ii) fuse the encoder output with attribute features, which increases the weights from attributes; (iii) propose a message-passing mechanism to enhance the information transfer between objects, relations and attributes. We conduct extensive experiments on datasets including GQA, VG, motif-VG and achieve new state of the art.

Index Terms— Scene Graph, Visual Question Answer, Graph Neural Network

1. INTRODUCTION

VQA tasks require a model to answer a free-form natural language question using visual information from an image. Scene graph (SG) reasoning is an essential instance of VQA tasks [1]. The model extracts objects' names, attributes, and relations from the input images and organizes them into a graph representation to generate the scene graph.

SG representation modeling displays several virtues over classical VQA techniques since the features in SG are presented in plain and free text form [2] and the graph struc-



Fig. 1. (a) Two key issues of traditional scene graph based models that we address: false attribute selection (top: select attribute "black" instead of attribute "off") and missing relation (bottom: missing "Behind") (b) Overview of our DM-GNN model. F^N and F^E are the object feature map and the relation feature map. F is the full-scale feature map.

tures of SG have better interpretability [3]. In this contribution, two reasoning methods on scene graphs are proposed: (i) consider scene graphs as probabilistic graphs and iteratively update nodes' probabilities using soft instructions extracted from questions [4]; (ii) apply Graph Neural Network (GNN) into scene graphs [5, 6] to learn joint representations of nodes and their relations, and then feed these representations into a predictor to get the answer. Scene graph reasoning frameworks are useful in VQA [7]. However, there still remain imperfections dealing with complex reasoning questions.

First, existing models tend to predict wrong answers for complex reasoning questions with attributes. Consider the

^{*}Corresponding Author

[†] Work was done during first author's internship at Baidu.

This work is supported in part by the Nature Science Foundation of China (No.61972217, No.62081360152), Natural Science Foundation of Guangdong Province in China (No.2019B1515120049, 2020B1111340056).



Fig. 2. Model structure of the Dual Message-passing enhanced Graph Neural Networks. MP stands for the message-passing module. Images are transformed into scene graphs by the scene graph generator. The object-significant form and relation-significant form of the scene graph are injected into the object encoder and the relation encoder. Nodes' representations are generated from the sum of the MP modules. The representations are then fused with question representations to predict answers.

"why" question in Fig. 1(a) as an example, **false attribute selection** occurs because the model cannot associate "off" relation with "dark monitor" object. The attribute selections require comprehensive supervision from objects, relations, and attributes, but existing methods focus too much on objects while ignoring attributes. Generally, information from objects and relations connected to them are reconstructed into object features in GNN-based methods [8]. However, these encoding methods lack information from objects' attributes. NSM methods [4] use soft instructions to update answer possibilities, but they treat attributes as secondary information.

Second, existing approaches answer poorly for complex questions that require information about relations. For instance, for localization questions, as in the "where" type of questions, in Fig. 1(a), we observe that **missing relation** occurs because the model cannot capture the relation information "Behind". Existing models have a strong bias towards object features, while considering relation as references. The unbalanced focus on objects and relations makes the models fail to learn discriminative representations for relations.

To improve the balance of all kinds of information in scene grpahs, we propose the Dual Message-passing enhanced Graph Neural Network (DM-GNN) for VQA, introducing a novel scene graph reasoning model that extracts balanced feature maps from objects, attributes, and relations information in scene graphs. Concretely, as shown in Fig. 1(b), our DM-GNN model is composed of a scene graph generator, a question encoder, dual graph encoders, and a fusion module. Besides, to balance the importance of objects and relations, we transform scene graphs into a relation-significant modality, where nodes represent relations and edges represent objects, and an object-significant modality, in which nodes represent objects and edges represent relations. After receiving scene graphs in two modalities, dual graph encoders can produce feature maps focusing on relations and objects.

Furthermore, to enhance the information transfer between

objects, relations and attributes, we modify the gated graph neural network (GGNN) structure in our DM-GNN by adding the message-passing module. It is a bidirectional GRU that guides the internal information flow. The encoder captures information from nodes, edges, and adjacent nodes that connect to them. In Fig. 1(b), the output feature map of the encoder passes through the fusion module, where the attribute features are explicitly modeled into the feature map to increase the information weight from attributes. Then the feature map passes through multi-head attention layers using question features extracted from the question encoder. Hence, the model dynamically focuses on the critical parts of the questions and uses the most similar part of the scene graph as the most adequate answer. Our main contributions are as follows:

- We analyse that existing models answer imperfectly for complex reasoning questions with attributes or relations due to the unbalance focus on three information types in scene graphs, which contain objects, relations and attributes.
- We propose a novel DM-GNN model containing a dual encoder structure and a message-passing module. Our model can obtain a balanced representation by properly encoding multi-scale scene graph information.
- Experimental results on various datasets show that DM-GNN effectively improves the reasoning accuracy on semantically complicated questions.

2. RELATED WORK

Visual Question Answering. Most VQA approaches use sequential models [9] to encode questions and CNN-based pretrained models [10] to encode images. Then they use attention methods [11, 12] to fuse features from images and questions. Transformer models [22] achieve outstanding performances on VQA tasks, yet they are heavy to train and hard to explain [23]. Instead, the scene graph model stands for an alternative that is more lightly and explainable.

Scene Graph Generation and Reasoning. Scene graph generation (SGG) methods [21] use object detection methods to extract region proposals from images. Scene graph can promote explainable reasoning for downstream multimodal tasks such as VQA [3]. In typical scene graph reasoning models, NSM [4] performs sequential reasoning over the scene graph by iteratively traversing its nodes. Other models [5, 6, 24] use GGNN [25] based model to encode scene graphs. However, previous works are hard to fully utilize the attribute information and learn the comprehensive representation of SG.

Graph Neural Network. GNN [14] is designed to infer on data described by graphs. [15, 16] apply GNN-based models on knowledge graphs, which are similar to scene graphs. However, existing GNN-based models cannot effectively process graphs with node attributes and complicated labels. Our DM-GNN model can learn a comprehensive and balanced representation using full-scale scene graph information from objects, attributes, and relations to overcome these problems.

3. DM-GNN METHODOLOGY

Our proposed architecture is illustrated in Fig. 2. We use the scene graph generator from [19]. In the question encoder, semantic questions are first projected into an embedding space using GLOVE pretrained word embedding model [17]. Then we use long short-term memory (LSTM) networks to generate questions representation $q \in \mathbb{R}^{dim}$, where dim is the dimension of the question representation. We introduce our dual graph encoders and fusion module in following subsections.

3.1. Object/Relation-Significant Graph

We organize scene graphs into object-significant graphs and relation-significant graphs.

Object-Significant Graph. We define the object significant graph as G_{obj} , where each node represents an object in the image and each edge represents a relation between two objects. Define *N* as the node set and *E* as the edge set. For $n_i, n_j \in N, e_k \in E, < n_i - e_k - n_j >$ denotes the relation tuple that represents relation e_k from object n_i to object n_j .

Relation-Significant Graph. We define relation significant modality as G_{rel} , where each node represents a relation between objects in the image and each edge represents an object, which is completely opposed to the object-significant modality. For $e_i, e_j \in E$, $n_k \in N$, $\langle e_i - n_k - e_j \rangle$ represents the relations e_i and e_j have a shared object n_k .

Attribute types. Define *L* as attribute types. For each node $n_i \in N$, we define a set of L + 1 property variables $\{n_i^l\}_{l=0}^L$, where n_i^0 represents n_i 's name embedding and n_i^l represents the embedding of node n_i 's l^{th} attribute.



Fig. 3. Overview of the pipeline at timestep t. (a) The message-passing module generates MP_{in} and MP_{out} as incident and output information gains. (b) The propagator uses information gains and h_i^{t-1} to update h_i^t .

3.2. Dual Encoders

We apply two GGNN-based encoders for object significant graph and relation significant graph. The encoder for object graph focuses on object features and the encoder for relation graph focuses on relation features. The dual structure can balance the importance of relations and objects.

Prior to encoding, every input scene graph is transformed into an information tuple (N, E, A_{in}, A_{out}) : N and E are collections of node embeddings and edge embeddings. A_{in} and A_{out} are the adjacency matrix of incident and output edges.

Let h_i^t is the hidden state of node n_i in the encoder at timestep t, then at t = 0, we initialize h_i^0 as the GLOVE embedding of n_i with zero padding:

Message-passing Module. To enhance the information transfer from edges and adjacent nodes to the updating nodes, we use the message-passing module (MP) in Fig. 3(a). MP module comes as a replacement of the fully-connected layers from the original GNN model. Consider a tuple $< n_i, e_k, n_j >$ as the processing sample of the message-passing module. The embedding state, noted e_k , of the edge e_k and neighbor node n_j 's hidden state h_j are injected into a bidirectional GRU network as input sequence while the node n_i 's hidden state h_i is injected as the GRU's initial hidden state. The output of the GRU represents the updating information for hidden state h_i , which corresponds to the key information from edge e_k and node n_j that is related to node n_i . The sum of every GRU output is n_i 's total information gain from n_i 's adjacent nodes and edges. We detail the message-passing

module formula as follows:

$$MP_i(A_{in}) = \sum_{k,j}^{\langle n_i, e_k, n_j \rangle \in A_{in}} \operatorname{GRU}([e_k, h_j], h_i), \quad (1)$$

$$MP_i(A_{out}) = \sum_{k,j}^{\langle n_j, e_k, n_i \rangle \in A_{out}} \operatorname{GRU}([e_k, h_j], h_i), \quad (2)$$

where $MP_i(A_{in})$ is n_i 's incident information gain, and $MP_i(A_{out})$ is n_i 's output information gain.

Propagation Module. In Fig. 3(b), at timestep t, the hidden states of all nodes are updated by the following gated propagator module:

$$k_i^t = [MP_i^t(A_{in}), MP_i^t(A_{out})], \qquad (3)$$

where k_i^t is the node n_i 's representation from all its incident edges, output edges and adjacent nodes.

Then, we incorporate information from adjacent nodes and from the previous timestep leading to an update of each node's hidden state:

$$c_i^t = [h_i^{(t-1)}, k_i^{(t-1)}]W + b, \qquad (4)$$

$$z_i^t = \sigma(U^z c_i^t) \,, \tag{5}$$

$$r_i^t = \sigma(U^r c_i^t) \,, \tag{6}$$

where W, U^z and U^r are referred to as the trainable weight matrices. At timestep t, z_i^t and r_i^t are the update and reset gates, respectively. Then we have:

$$\tilde{h}_{i}^{t} = \tanh(U_{1}k_{i}^{(t-1)} + U_{2}(r_{i}^{t} \odot h_{i}^{(t-1)})), \qquad (7)$$

$$h_{i}^{t} = (1 - z_{i}^{t}) \odot h_{i}^{(t-1)} + z_{i}^{t} \odot \tilde{h}_{i}^{t}.$$
(8)

Here, U_1 and U_2 denote the trainable parameters of the linear layers, the operator \odot is the element-wise multiplication. σ is the ReLU function. After T steps, the encoder generates the final hidden state map G of the graph. Finally, we compute the graph embedding $g_i \in G$ for node n_i as follows:

$$g_i = \sigma(f(h_i^T, n_i)), \qquad (9)$$

where $f(h_i^T, n_i)$ is multi-layer perceptron (MLP) which receives the concatenation of h_i^T and n_i .

3.3. Fusion Module and Answer Predictor

Once the dual encoders, embedded in our model, output the node and relation features, we first fuse the attributes into feature maps. For node feature map G^N and relation feature map G^E , the fusion feature map F^N and F^E are defined as

$$F_i^N = \begin{cases} [g_i^N, n_i^0] \\ \cdots, \\ [g_i^N, n_i^L] \end{cases}, \ F_j^E = [g_j^E, e_j], \ F = [F^N, F^E], \quad (10)$$

where F_i^N indicates the fusion features of node *i* and g_i^N is node *i*'s representation from the encoder. n_i^l is the attribute embedding of node *i*. F_j^E corresponds to the fusion feature of edge *j*. g_j^E is edge *j*'s representation from the encoder. e_j is *j*-th edge original embedding. The full-scale feature map, noted *F*, is obtained by concatenating F^N and F^E .

Then, the question embedding q and the full-scale feature map F are fed into a multi-head attention layer. The reasoning vector, noted r, and which stems from the graph and the question, is computed using a weighted sum of the feature map using the scores output from the attention layer,

$$r = \operatorname{Attention}(F, q)$$
. (11)

In answer predictor module, we adopt a two-layer MLP noted by $f(\cdot)$. This MLP can be viewed as a classifier over the set of candidate answers. The input of the answer predictor is the concatenation vector (q, r). Such a classifier has been applied in many VQA models [4, 11]. The answer \hat{a} reads:

$$\hat{a} = \arg\max(\operatorname{softmax}(f((q, r))))).$$
(12)

4. EXPERIMENTS

4.1. Empirical Results

Results on VG dataset. Table 1 reports the results on the test sets of the VG ground truth dataset and the motif-VG dataset. Compared to the baseline models, we can observe that our DM-GNN model outperforms the others at **3%-4%**. In addition, we provide detailed results on the VG dataset and motif-VG dataset with different question types. Compared to the other scene graph based VQA models, our model performs well in "what", "where", "who" and "why" types. On the VG dataset, our model has **6.5%** accuracy improvement in "**why**" type questions, which highly requires VQA models' ability to jointly exploit objects, relations and attributes.

Results on GQA dataset. We report in Table 2 the detailed results on the test sets of the GQA dataset. Our DM-GNN model outperforms baselines. We also evaluate our model and other baselines across GQA dataset's various metrics, where "Binary" and "Open" stand for binary-answer and open domain questions. "Distribution" corresponds to the distance between prediction distribution and standard answer distribution. For open domain questions which are difficult for reasoning, our model outperforms the others by 9.63%. For distribution metric, our model also achieves 2nd score. However, the "Binary" question is challenging for DM-GNN although it achieves SOTA performances. Specifically, the advantage of DM-GNN is to correctly locate key objects, relations and attributes. That is why it works well for the open domain ("what, why, how") questions. But for "yes/no" questions, whose answers are not explicit in scene graphs, DM-GNN is easy to locate but hard to correctly answer.

Dataset	VG-GroundTruth					Motif-VG				
Question type	What	Where	Who	Why	Overall	What	Where	Who	Why	Overall
Percentage	(54%)	(17%)	(5%)	(3%)	(100%)	(54%)	(17%)	(5%)	(3%)	(100%)
NSM [4]	33.1	51.0	49.8	12.3	45.1	31.8	53.1	47.6	10.9	43.1
F-GN [3]	60.9	62.0	63.3	50.9	60.1	58.7	60.4	61.8	49.0	60.0
U-GN [3]	61.6	62.4	63.9	50.3	60.5	59.4	60.3	66.6	48.1	60.5
FSTT [5]	65.5	70.1	68.3	91.5	65.6	48.8	49.2	40.6	70.3	48.1
ReGAT [6]	72.1	64.4	72.7	92.3	71.2	75.4	57.6	69.1	91.8	69.9
DM-GNN (ours)	75.9	73.1	82.6	98.8	75.4	79.4	62.7	72.8	96.1	72.9

Table 1. Performance on different question types of VG dataset.

Models	Binary↑	Open↑	Validity	Distribution ↓	Acc.↑	Models	Acc.
Human	91.20	87.40	98.90	-	89.30	Base	35.4
BottomUp	66.64	34.83	96.18	5.98	49.74	Base-Obj	35.4
MAC	71.23	38.91	96.16	5.34	54.06	+MP	39.3 _(+3.9)
SK T-Brain	77.42	43.10	96.26	7.54	59.19	+Dual	$67.9_{(+32.7)}$
PVR	77.69	43.01	96.45	5.80	59.27	Base-Rel	35.2
GRN	77.53	43.35	96.18	6.06	59.37	+MP	$38.8_{(+3.6)}$
Dream	77.84	43.72	96.38	8.40	59.72	+Dual	$67.9_{(+32.5)}$
LXRT	77.76	44.97	96.30	8.31	60.34	DM-GNN(ours)	75.2
NSM	78.94	49.25	96.41	3.71	63.17	+ (w/o <i>attr</i>)	$71.6_{(-3.6)}$
ReGAT	83.57	62.58	92.70	9.32	70.50	+ (w/o <i>rela</i>)	$74.5_{(-0.7)}$
DM-GNN(ours)	69.79	72.21	93.80	3.78	71.21	+ (w/o <i>QF</i>)	$54.9_{(-20.3)}$

Table 2. Performance on the GQA dataset.

Table 3. Ablation Study on VG.

4.2. Ablation Study

We compare several ablated forms of DM-GNN with our complete model on the VG dataset. The accuracy for each variant of DM-GNN are reported in Table 3. We use the raw GGNN network as the *Base* model. The *Base-Obj* and *Base-Rel* model represent the original GGNN network processing *object-significant* graph and *relation-significant* graph. The models with +*MP* contain the message-passing module. The models with +*Dual* apply the dual encoder structure.

Effect of dual encoder structure. We first validate the efficacy of applying dual structure to balance the importance of relations and objects by splitting our DM-GNN into two single models (*Base-Obj* and *Base-Rel*). Both single models perform poorly at $35.3\% \pm 0.1\%$. This also shows that both relations and objects are vital to VQA performance. Absence of any of those modules leads to severe accuracy recession. Adding dual encoder structure (+*Dual*) leads to an empirical gain of **32.6%** accuracy upward, which shows that the dual structure is significant in balancing relations and objects.

Effect of the message-passing module. We validate the effectiveness of applying message-passing structure to learn a more comprehensive representation for scene graphs than the raw GGNN structure. Comparing with the *Base-Obj* and +*MP* model, we note that after adding the message-passing structure, there is an improvement of 3.9%. Comparing with the *Base-Rel* and +*MP* model, we observe an improvement of 3.6%. The +*Dual* model is DM-GNN without message-passing module. Comparing with DM-GNN, there

is a 7.3% improvement after adding the message-passing module. These empirical results show that message-passing structure can successfully improve the representation quality of scene graphs.

Effect of explicit modeling. The *w/o attr* model and the *w/o rela* model remove the explicit attribute modeling and relation modeling part. Comparing *w/o attr*, *w/o rela* with DM-GNN, removing attribute modeling has a 3.6% decrease in accuracy and removing relation modeling has 0.7% decrease.

4.3. Visualization

Fig. 4 (a) shows visualization on "what" question type. Three "what" question examples aimed at retrieving either object, relation or attribute information. Comparing row 1, row 2 with row 3, *Obj* and *Rel* models have strong attention bias toward objects and relations, while their combination Obj+Rel, balances the attention on both sides and captures correct answers. Comparing row 3 with row 4, the message-passing module increases the score of correct answers.

Fig. 4 (b) shows the visualization on "why", which need models to jointly exploit objects, relations and attributes to infer answers. With dual encoders and the message-passing module, our DM-GNN achieves **96.1%** on "why" questions.

5. CONCLUSION

We propose DM-GNN, which encodes each scene graph into feature representations via an object encoder and a relation encoder generating a balanced and full-scale feature map using objects, attributes, and relations information, and demon-



Fig. 4. Examples of generated answers and top attention scores. Answers in red means wrong and green means right. (a) and (b) show the visualization of "what" and "why" question types. From (a), we observe that our approach can correctly select answer from objects, relations and attributes. From (b), we note that our model can handle comprehensive reasoning questions.

strate our model can effectively boost performances on GQA, VG and Motif-VG datasets .

6. REFERENCES

- M. Hildebrandt, H. Li, R. Koner, V. Tresp, and S. Günnemann, "Scene graph reasoning for visual question answering," *CoRR*, 2020.
- [2] V. Damodaran, S. Chakravarthy, A. Kumar, A. Umapathy, T. Mitamura, Y. Nakashima, N. Garcia, and C. Chu, "Understanding the role of scene graphs in visual question answering," *CoRR*, vol. abs/2101.05479, 2021.
- [3] C. Zhang, W. Chao, and D. Xuan, "An empirical study on leveraging scene graphs for visual question answering," in *BMVC 2019*.
- [4] D.A. Hudson and C.D. Manning, "Learning by abstraction: The neural state machine," in *NeurIPS 2019*.
- [5] Aj. Singh, An. Mishra, Sh. Shekhar, and An. Chakraborty, "From strings to things: Knowledgeenabled vqa model that can read and reason," .
- [6] L. Li, Z. Gan, Y. Cheng, and J. Liu, "Relation-aware graph attention network for visual question answering," in *ICCV 2019*.
- [7] Z. Yang, Z. Qin, J. Yu, and T. Wan, "Prior visual relationship reasoning for visual question answering," in *ICIP 2020*.
- [8] H. Xu, C. Jiang, X. Liang, and Z. Li, "Spatial-aware graph relation network for large-scale object detection," in *CVPR*, 2019.
- [9] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in NAACL-HLT 2019.

- [10] Badri N. Patro and Vinay P. Namboodiri, "Differential attention for visual question answering," in 2018 CVPR.
- [11] J. Lu, J. Yang, D. Batra, and D. Parikh, "Hierarchical question-image co-attention for visual question answering," in *NIPS 2016*.
- [12] D.A. Hudson and C.D. Manning, "Compositional attention networks for machine reasoning," in *ICLR 2018*.
- [13] Q. Cao, X. Liang, B. Li, and L. Lin, "Interpretable visual question answering by reasoning on dependency trees," *IEEE TPAMI.*, 2021.
- [14] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," .
- [15] Y. Wang, J. Guo, W. Che, and T. Liu, "Transition-based chinese semantic dependency graph parsing," in *NLP-NABD 2016*.
- [16] Y. Wang, W. Che, J. Guo, and T. Liu, "A neural transition-based approach for semantic dependency graph parsing," in *AAAI 2018*.
- [17] J. P., R. Socher, and C. Manning, "GloVe: Global vectors for word representation," in *EMNLP 2014*.
- [18] D.A. Hudson and C.D. Manning, "GQA: A new dataset for real-world visual reasoning and compositional question answering," in *CVPR 2019*.
- [19] K. Tang, "Sgg codebase in pytorch," 2020.
- [20] G. Yin, L. Sheng, B. Liu, N. Yu, X. Wang, J. Shao, and Chen C., "Zoom-net: Mining deep feature interactions for visual relationship recognition," in *ECCV 2018*.
- [21] K. Tang, Y. Niu, J. Huang, J. Shi, and H. Zhang, "Unbiased scene graph generation from biased training," in *CVPR 2020*.
- [22] G. Li, N. Duan, Y. Fang, Ming G, and Daxin J, "Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training," in *AAAI 2020*.

- [23] X. Han, Z. Zhang, N. Ding, J. Wen, J. Yuan, W. Zhao, and J. Zhu, "Pre-trained models: Past, present and future," *CoRR*, 2021.
- [24] W. Liange, Y. Jiang, and Z. Liu, "Graphvqa: Languageguided graph neural networks for scene graph question answering," *NAACL 2021*.
- [25] Y. Li, D. Tarlow, M. Brockschmidt, and R.S. Zemel, "Gated graph sequence neural networks," in *ICLR 2016*.
- [26] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L. Li, D.A. Shamma, M.S. Bernstein, and Li Fei-Fei, "Visual genome: Connecting language and vision using crowdsourced dense image annotations," *Int. J. Comput. Vis.*, 2017.