

# SANVis: Visual Analytics for Understanding Self-Attention Networks

Cheonbok Park\* Korea University  
 Bum Chul Kwon§ IBM Research  
 Inyoup Na\* Korea University  
 Jian Zhao¶ University of Waterloo  
 Yongjang Jo† Korea University  
 Hyungjong Noh|| NCSOFT Co., LTD.  
 Sungbok Shin‡ University of Maryland  
 Yeonsoo Lee|| NCSOFT Co., LTD.  
 Jaehyo Yoo\* Korea University  
 Jaegul Choo\* Korea University



Figure 1: Overview of SANVis. (A) The network view displays multiple attention patterns for each layer according to three type of visualization options: (A-1) the attention piling option, (A-2) the Sankey diagram option, and (A-3) the small multiples option. (A-4) The bar chart shows the average attention weights for all heads (each colored with its corresponding hue) per each layer. (B) The HeadLens view helps the user analyze what the attention head learned by showing representative words and by providing statistical information of part-of-speech tags and positions.

## ABSTRACT

Attention networks, a deep neural network architecture inspired by humans’ attention mechanism, have seen significant success in image captioning, machine translation, and many other applications. Recently, they have been further evolved into an advanced approach called multi-head self-attention networks, which can encode a set of input vectors, e.g., word vectors in a sentence, into another set of vectors. Such encoding aims at simultaneously capturing diverse syntactic and semantic features within a set, each of which corresponds to a particular attention head, forming altogether multi-head attention. Meanwhile, the increased model complexity prevents users from easily understanding and manipulating the inner workings of models. To tackle the challenges, we present a visual analytics system called SANVis, which helps users understand the behaviors and the characteristics of multi-head self-attention networks. Using a state-of-the-art self-attention model called Transformer, we demonstrate usage scenarios of SANVis in machine translation tasks. Our system is available at <http://short.sanvis.org>.

**Index Terms:** Deep neural networks, visual analytics, natural language processing, interpretability, self-attention networks

## 1 INTRODUCTION

Attention-based deep neural networks, inspired by humans’ attention mechanism, are widely used for sequence-to-sequence modeling, e.g., neural machine translation. The attention module allows the model to dynamically utilize different parts of the input sequence, which leads to state-of-the-art performances in natural language processing (NLP) tasks [4, 14, 32].

However, conventional approaches using recurrent neural networks (RNNs) had limitations that (1) they utilize only a single attention module that can capture only a particular characteristic of a given input and that (2) they cannot properly capture long-range dependencies due to the loss in memory content over multiple time steps.

To address these limitations, Vaswani et al. [26] recently proposed multi-head self-attention networks (in short, self-attention networks), which replace an RNN-based sequence encoding module with a sophisticated attention module. This module is composed of multiple different attention heads, each of which captures its own syntactic and/or semantic features within a set. Owing to these advantages, self-attention-based models have achieved state-of-the-art performances in machine translation, and it has been further extended in other NLP tasks [8, 20] and computer vision domains [31, 35].

However, their highly sophisticated model architecture prevents users from deeply understanding and interacting with them. In response, this paper presents a visual analytics system for self-attention networks, called SANVis, as well as its comprehensive usage scenarios using widely-used networks called Transformer [26].

\* {cb\_park, windy9898, vkfwlsdhs, jchoo} @korea.ac.kr

† jyj3312@gmail.com

‡ sbshin90@cs.umd.edu

§ bumchul.kwon@us.ibm.com

¶ jianzhao@uwaterloo.ca; work was completed while at FXPAL.

|| {nohhj0209, yeonsoo} @ncsoft.com

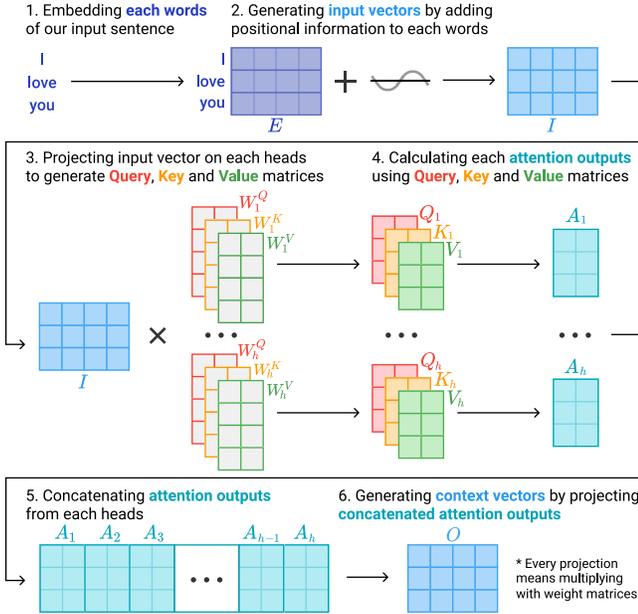


Figure 2: How a multi-head self-attention module works. Steps 1 and 2 correspond to the embedding layer, while Steps 3 to 6 correspond to a single-layer multi-head self-attention module.

## 2 RELATED WORK

We discuss related work from two perspectives: (1) visual analytic approaches for interpreting and interacting with various deep neural networks and (2) interpretation and analysis of self-attention networks mainly in NLP domains.

Regarding the former, various visual analytic approaches have been proposed for convolutional neural networks (CNNs) mainly computer vision domains [2, 6, 12, 13, 19, 34] and RNNs in NLP domains [5, 11, 17, 23, 24]. Visual analytic approaches have also been integrated with other advanced neural network architectures, such as generative adversarial networks [9, 30], deep reinforcement learning [29]. Among them, Strobel et al. [22] developed a visual analytic system for RNN-based attention models, mainly for the exploration and understanding of sequence-to-sequence modeling tasks. However, despite the success of multi-head self-attention networks, such as BERT [8] and XLNet [33], visual analytic approaches for these advanced attention networks have not existed before.

In NLP domains, recent studies [7, 8, 27] have analyzed diverse behaviors of different attention heads in a self-attention model and have drawn linguistic interpretations as to what kind of syntactic and/or semantic features each attention head captures. Another line of research [25, 28] have attempted to leverage insights obtained from such in-depth analysis to improve the prediction accuracy and computational efficiency by removing unnecessary heads and refining them. However, these approaches have not properly utilized the potential of interactive visualization, so in this respect, our work is one of the first sophisticated visual analytics systems for self-attention networks.

## 3 SELF-ATTENTION NETWORKS

This section briefly reviews the self-attention module originally proposed in Transformer [26]. Transformer adopts an encoder-decoder architecture to solve sequence-to-sequence learning tasks, e.g., neural machine translation, which converts a sentence in one language into that in another language. It converts a sequence of words in one domain into that in another domain. For example, for machine translation tasks, it translates a sentence in one language into that in another language. In this process, the encoder of Transformer converts input words (e.g., English words) to internal, hidden-state

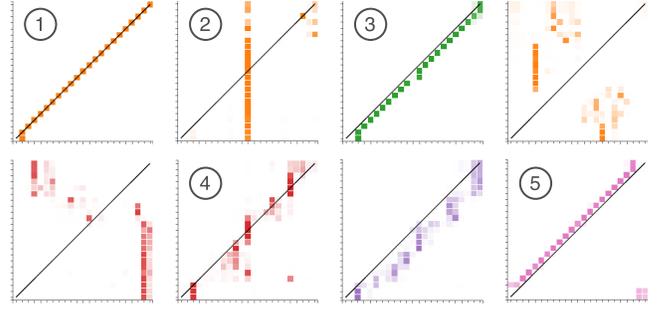


Figure 3: Diverse attention patterns found in the encoder of Transformer. Some attention heads show diagonal patterns indicating that a query word attends to itself (1) or its immediate previous (5) or next word (3). Some other attention heads attend to a common, single word (2). In other attention heads, each group of consecutive words attends commonly to a single word within that group (4).

vectors, and the decoder turns the vectors into a sequence of output words (e.g., French words).

Each encoder and decoder respectively consists of multiple layers of computing functions inside. Furthermore, each layer in the encoder includes two sequential sub-layers, which are a multi-head self-attention and a position-wise feed-forward network. In addition to the multi-layer architecture of the encoder, the decoder has an additional attention layer, which is called as an encoder-decoder attention and helps the model to give attention to the encoders' internal states. Each layer of both encoder and decoder also consists of skip-connection and layer normalization in their computation pipeline. Overall encoder and decoder architecture are the stacks of  $L$  identical encoder layers or decoder layers, including an embedding layer.

We summarize the computation process with mathematical notations, so readers are advised to read the remaining section for details: Let us denote  $d_{model}$  as the size of hidden state vector and  $h$  as the number of heads in multi-head self-attention. Each dimension of query, key, and value vector is  $d_q = d_k = d_v = d_{model}/h$ .

The embedding layer transforms the input token  $x_i$  to its embedding space  $e_i$  using a word embedding and adds the position information for each input token using sinusoidal functions (see Steps 1 and 2 in Figure 2), where  $x_i$  is the  $i$ -th input token in  $X = [x_1, \dots, x_T]$ .

At each attention head, we transform encoded word vectors into three matrices of a query, a key, and a value,  $Q \in R^{T \times d_q}$ ,  $K \in R^{T \times d_k}$ , and  $V \in R^{T \times d_v}$ , respectively, for  $h$  times, which in turn generates  $h \times 3$  matrices, using the linear transformation and computes the attention-weighted combinations of value vectors as

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d_{model}}}\right)V$$

$$\text{MultiHeadAttention} = \text{Concat}(\text{head}_1, \dots, \text{head}_h)W^O \quad (1)$$

where  $\text{head}_i = \text{Attention}\left(QW_i^Q, KW_i^K, VW_i^V\right)$ , and  $W_i^Q$ ,  $W_i^K$  and  $W_i^V$  indicate the linear transformation matrices at the  $i$ -th head. In multi-head self-attention, which consists of  $h$  parallel attention heads, transformation matrices of each head are randomly initialized, and then each set is used to project input vectors onto a different representation subspace. For this reason, every attention head is allowed to have different attention shapes and patterns. This characteristic encourages each head differently to attend adjacent words or linguistics relation words.

In the encoder layer, source words (input words to the encoder) work as the input to the query, key, and value transformations at the  $i$ -th head. In the decoder layer, the input can vary by attention types. While the decoders' self-attention takes target words (output words of the decoder) as its input, the encoder-decoder attention has target words as input to a query transformation but source words as the input to a key and a value transformation.

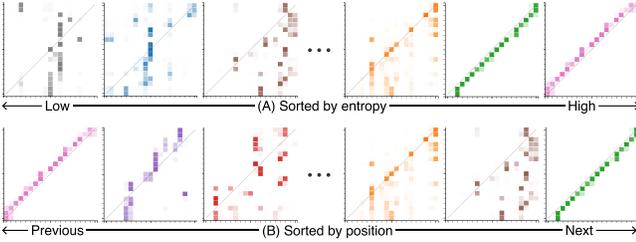


Figure 4: Attention sorting result. The user can sort a set of multiple attention patterns with respect to different criteria such as the entropy measure (A) and the relative positional offset from query words (B).

#### 4 DESIGN RATIONALE

We consider our design rationale of SANVis as follows:

- R1: Grasping the information flow across multiple layers.**
- R2: Identifying and making sense of attention patterns of each attention head.**
- R3: Visualizing syntactic and semantic information to allow of exploring the attention head in their query and key vectors.**

#### 5 SANVis

We present SANVis,<sup>1</sup> a visual analytics system for deep understanding of self-attention models, as shown in Figure 1. SANVis is composed of the network view and the HeadLens view. (1) The network view allows the user to understand the overall information flow through our visualization across the multiple layers (T1). Moreover, this view provides additional visualization options that assist the user in distinguishing distinct patterns from multiple attention patterns within a layer (T2). (2) The HeadLens view reveals the characteristics of the query and the key vectors and their relationship of a particular head (T3).

##### 5.1 Network View

Network view mainly visualizes the overview of attention patterns across multiple layers using the Sankey diagram (T1). Additionally, this view supports ‘piling’ and ‘sorting’ capabilities to understand common as well as distinct attention patterns among multiple attention heads (T2). For example, one can replace the Sankey diagram with a multiple heatmap view, where multiple heatmaps corresponding to different heads can be sorted by several different criteria (Figure, 1 (A-3)). The attention piling view aggregates multiple attention patterns into a small number of clusters (Figure 1(A-1)).

As shown in Figure 1(A), a set of words are sequentially aligned vertically in each layer, and represented the histogram according to attention weights from multiple heads. In Figure 1(A-4), each bar corresponds to a particular head within the layer where its height represents the total amount of attention weights assigned to those words by a specific head. As with Figure 1(A-2), if the fourth head in the layer attended to the word ‘planet’ more highly than others, the fourth bar would be higher than the others. In this manner, the user easily recognizes which heads highly attend those words based on the height of histogram bars. Furthermore, when the user moves the mouse over the particular color bar, we show an attention heatmap of the corresponding head in that layer.

**Sankey diagram.** As shown in Figure 1(A-2), the edge weight between them represents the average attention weight across multiple heads within a particular layer. In this figure, we can see the strong link between ‘planet’ in layer 2 to the preposition words(‘on’), the same word and that article(‘the’) in layer 3. It means a significant amount of information of ‘planet’ in layer 2 is conveyed to encode each word of a phrase (“on the planet”) in layer 3. This pattern shows the model captures the context meaning of the word, which is

<sup>1</sup>Our demo is available at <http://short.sanvis.org>.

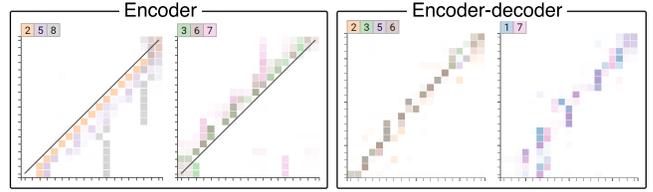


Figure 5: Attention piling example in the encoder layer and encoder-decoder layer. In the encoder-decoder example, piling results do not have a gray diagonal line because of the difference between the count of query words and key words.

defined as the linguistic phrase in the given sentence, for improving the quality of translation.

**Attention Sorting.** Figure 3 shows various attention patterns between query (y-axis) and key (x-axis) words for the different attention head in different layers, where a gray diagonal line indicates the position of attending itself for detecting attention patterns.

We focus on reducing the users’ efforts to find the distinguish attention patterns by using sorting algorithms, which is based on relative positional information and the entropy value in the attention (Figure 3). Relative positional information, such as whether the attention goes mainly toward the left, right, or the current location, as well as the column-wise mean entropy value of the attention matrix, allow the users to detect these patterns easily.

Figure 4 shows the sorted results of attention patterns based on our position or entropy sorting algorithms. When sorted by position, a number of attention is unambiguous that attention that inclines towards the past words are placed near the control panel at the top while those that lean towards the future words are placed relatively close to the bottom. When sorted by entropy, the uppermost attention has the lowest entropy and exhibits bar-shaped attention, which numerous query words attend the same word. At the bottom, the user can find that no more words focused on the same word.

**Attention Piling.** Inspired by the heatmap piling methods [3, 21], we applied this piling idea to summarize multiple attention patterns in a single layer, as shown in the encoder part of Figure 5. To this end, we compute the feature vector of each attention head and perform clustering to the form of piles (or clusters) of multiple attention patterns.

The feature vector of a particular attention on the attention head is defined as a flattened  $n^2$ -dimensional vector of its  $A_i \in R^{T \times T}$  attention matrix, where  $A_i$  is calculated from  $\text{Softmax}\left(\frac{QK^T}{\sqrt{d_{model}}}\right)$  on the  $i$ -th head, concatenated with additional three-dimensional vector of (1) the sum of the upper triangular part of the matrix, (2) that of the lower-triangular part, and (3) the sum of diagonal entries. This three-dimensional vector indicates the proportions how much attention is assigned to (1) the previous words of a query word, (2) its next words, (3) and the query itself, respectively.

Using these feature vectors of multiple attention heads within a single layer, we perform hierarchical clustering based on their Euclidean distances. In this manner, multiple attention patterns are grouped, forming an aggregated heatmap visualization per computed pile along with head indices belonging to each pile, as shown in Figure 5. It helps the user to easily find the similar patterns and distinct patterns in the same layer by adjusting Euclidean distance.

##### 5.2 HeadLens

To analyze a particular attention head, SANVis offers a novel view called the HeadLens, as shown in Figure 1(B). This view facilitates detailed analysis of the query and key representations of the selected attention head, such as which linguistic or positional feature they encoded (T3). This view opens when a user clicks a particular heatmap corresponding to an attention head in the network view.

The HeadLens is generated as follows. (1) It performs clustering on query and key vectors separately. (2) For each pair of a query and



- [3] B. Bach, N. H. Riche, T. Dwyer, T. M. Madhyastha, J.-D. Fekete, and T. J. Grabowski. Small multiples: Piling time to explore temporal patterns in dynamic networks. *Proceedings of EG/NGCT Conference on Visualization (EuroVis)*, 34:31–40, 2015.
- [4] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *Proc. the International Conference on Learning Representations (ICLR)*, 2014.
- [5] D. Cashman, G. Patterson, A. Mosca, N. Watts, S. Robinson, and R. Chang. RNNbow: Visualizing Learning Via Backpropagation Gradients in RNNs. *IEEE Computer Graphics and Applications*, 38(6):39–50, 2018.
- [6] S. Chung, C. Park, S. Suh, K. Kang, J. Choo, and B. C. Kwon. Revacnn: Steering convolutional neural network via real-time visual analytics. *Future of Interactive Learning Machines Workshop (FILM at NerIPS)*, 2016.
- [7] K. Clark, U. Khandelwal, O. Levy, and C. D. Manning. What does bert look at? an analysis of bert’s attention. *ArXiv*, abs/1906.04341, 2019.
- [8] J. Devlin, M. Chang, K. Lee, and K. Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *Proc. the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.
- [9] M. Kahng, N. Thorat, D. H. Chau, F. Viégas, and M. Wattenberg. GAN Lab: Understanding Complex Deep Generative Models using Interactive Visual Experimentation. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 25:310–320, 2019.
- [10] G. Klein, Y. Kim, Y. Deng, J. Senellart, and A. M. Rush. Opennmt: Open-source toolkit for neural machine translation. In *Proc. the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2017.
- [11] B. C. Kwon, M. Choi, J. T. Kim, E. Choi, Y. B. Kim, S. Kwon, J. Sun, and J. Choo. RetainVis: Visual Analytics with Interpretable and Interactive Recurrent Neural Networks on Electronic Medical Records. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):299–309, 2019.
- [12] D. Liu, W. Cui, K. Jin, Y. Guo, and H. Qu. DeepTracker: Visualizing the Training Process of Convolutional Neural Networks. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2019.
- [13] M. Liu, J. Shi, Y. Li, C. Li, J. Zhu, and S. Liu. Towards better analysis of deep convolutional neural networks. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 23:91–100, 2016.
- [14] T. Luong, H. Pham, and C. D. Manning. Effective approaches to attention-based neural machine translation. *Proc. the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1412–1421, 2015.
- [15] C. Manning, M. Surdeanu, J. Bauer, J. Finkel, S. Bethard, and D. McClosky. The stanford corenlp natural language processing toolkit. In *Proc. the Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pp. 55–60, 2014.
- [16] P. Michel and G. Neubig. Extreme adaptation for personalized neural machine translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 312–318. Association for Computational Linguistics, 2018.
- [17] Y. Ming, S. Cao, R. Zhang, Z. Li, Y. Chen, Y. Song, and H. Qu. Understanding hidden memories of recurrent neural networks. *IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 13–24, 2017.
- [18] R. Ostrovsky, Y. Rabani, L. J. Schulman, and C. Swamy. The effectiveness of lloyd-type methods for the k-means problem. *IEEE Symposium on Foundations of Computer Science (FOCS)*, (165-176).
- [19] N. Pezzotti, T. Höllt, J. V. Gemert, B. P. F. Lelieveldt, E. Eisemann, and A. Vilanova. Deepeyes: Progressive visual analytics for designing deep neural networks. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 24(1):98–108, 2018.
- [20] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. Language models are unsupervised multitask learners. 2019.
- [21] N. H. Riche, J.-D. Fekete, and M. J. McGuffin. Nodetrix: a hybrid visualization of social networks. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 13, 2007.
- [22] H. Strobel, S. Gehrmann, M. Behrisch, A. Perer, H. Pfister, and A. M. Rush. Seq2seq-vis: A visual debugging tool for sequence-to-sequence models. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 25:353–363, 2018.
- [23] H. Strobel, S. Gehrmann, M. Behrisch, A. Perer, H. Pfister, and A. M. Rush. SEQ2seq-VIS : A Visual Debugging Tool for Sequence-to-Sequence Models. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):353–363, 2019.
- [24] H. Strobel, S. Gehrmann, H. Pfister, and A. M. Rush. Lstmvis: A tool for visual analysis of hidden state dynamics in recurrent neural networks. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 24(1):667–676, 2018.
- [25] E. Strubell, P. Verga, D. Andor, D. Weiss, and A. McCallum. Linguistically-informed self-attention for semantic role labeling. *Proc. the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 5027–5038, 2018.
- [26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin. Attention is all you need. In *Proc. the Advances in Neural Information Processing Systems (NeurIPS)*, pp. 5998–6008. 2017.
- [27] J. Vig and Y. Belinkov. Analyzing the structure of attention in a transformer language model. *ArXiv*, abs/1906.04284, 2019.
- [28] E. Voita, D. Talbot, F. Moiseev, R. Sennrich, and I. Titov. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. *ArXiv*, abs/1905.09418, 2019.
- [29] J. Wang, L. Gou, H. Shen, and H. Yang. Dqnviz: A visual analytics approach to understand deep q-networks. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):288–298, 2019.
- [30] J. Wang, L. Gou, H. Yang, and H.-W. Shen. Ganviz: A visual analytics approach to understand the adversarial game. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 24:1905–1917, 2018.
- [31] X. Wang, R. B. Girshick, A. Gupta, and K. He. Non-local neural networks. *The Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7794–7803, 2018.
- [32] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. *Proc. the International Conference on Machine Learning (ICML)*, 37:2048–2057, 2015.
- [33] Z. Yang, Z. Dai, Y. Yang, J. G. Carbonell, R. Salakhutdinov, and Q. V. Le. Xlnet: Generalized autoregressive pretraining for language understanding. *ArXiv*, abs/1906.08237, 2019.
- [34] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *Proc. the European Conference on Computer Vision*, pp. 818–833. 2014.
- [35] H. Zhang, I. J. Goodfellow, D. N. Metaxas, and A. Odena. Self-attention generative adversarial networks. *ArXiv*, abs/1805.08318, 2018.