

Unsupervised Keyphrase Extraction by Jointly Modeling Local and Global Context

Xinnian Liang^{1*}, Shuangzhi Wu², Mu Li² and Zhoujun Li^{1†}

¹State Key Lab of Software Development Environment, Beihang University, Beijing, China

²Tencent Cloud Xiaowei, Beijing, China

{xnliang, lizj}@buaa.edu.cn; {frostwu, ethanlli}@tencent.com;

Abstract

Embedding based methods are widely used for unsupervised keyphrase extraction (UKE) tasks. Generally, these methods simply calculate similarities between phrase embeddings and document embedding, which is insufficient to capture different context for a more effective UKE model. In this paper, we propose a novel method for UKE, where local and global contexts are jointly modeled. From a global view, we calculate the similarity between a certain phrase and the whole document in the vector space as transitional embedding based models do. In terms of the local view, we first build a graph structure based on the document where phrases are regarded as vertices and the edges are similarities between vertices. Then, we proposed a new centrality computation method to capture local salient information based on the graph structure. Finally, we further combine the modeling of global and local context for ranking. We evaluate our models on three public benchmarks (Inspec, DUC 2001, SemEval 2010) and compare with existing state-of-the-art models. The results show that our model outperforms most models while generalizing better on input documents with different domains and length. Additional ablation study shows that both the local and global information is crucial for unsupervised keyphrase extraction tasks.

1 Introduction

Keyphrase extraction (KE) task aims to extract a set of words or phrases from a document that can represent the salient information of the document (Hasan and Ng, 2014). KE models can be divided into supervised and unsupervised. Supervised methods need large-scale annotated training data and always perform poorly when transferred to different domain or type datasets. Compared with the supervised method, the unsupervised method is more

universal and adaptive via extracting phrases based on information from input document itself. In this paper, we focus on the unsupervised keyphrase extraction (UKE) model.

UKE has been widely studied (Mihalcea, 2004; Wan and Xiao, 2008a; Bougouin et al., 2013; Boudin, 2018; Bennani-Smires et al., 2018; Sun et al., 2020) in the keyphrase extraction field. Recently, with the development of text representation, embedding-based models (Bennani-Smires et al., 2018; Sun et al., 2020) have achieved promising results and become the new state-of-the-art models. Usually, these methods compute phrase embeddings and document embedding with static word2vec models (e.g. GloVe (Pennington et al., 2014; Le and Mikolov, 2014; Pagliardini et al., 2018)) or dynamic pre-trained language models (e.g. BERT (Devlin et al., 2019)). Then, they rank candidate phrases by computing the similarity between phrases and the whole document in the vector space. Though, these methods performed better than traditional methods (Mihalcea, 2004; Wan and Xiao, 2008a; Bougouin et al., 2013), the simple similarity between phrase and document is insufficient to capture different kinds of context and limits in performance.

Figure 1 shows an intuitive explanation for the importance of context modeling. The nodes are candidate phrase embeddings, the star is the document embedding. Each black circle represents one local context. Nodes in the same black circle mean that these candidate phrases are all related to one vital local information (e.g. one topic/aspect of the document). Nodes in the red circle mean that these candidate phrases are similar with the document semantics. If only model the global context via computing similarity between candidate phrases and the document, the model will tend to select red nodes, which will ignore local salient information in three clusters. In order to get the keywords accurately, we should take the local con-

*Contribution during internship at Tencent Inc.

†Corresponding Author




Figure 1: Visualization of embedding space. Nodes refer to candidate phrase representation and star is document representation. Black circles mean clusters which contain local salient information. Red circle means global similarity phrases.

text (black circles) and global context (red circle) into consideration.

To obtain information from context adequately, in this paper, we proposed a novel method which jointly models the local and global context of the input document. Specifically, we calculate the similarity between candidate phrases and the whole document for modeling global context. For local context modeling, we first build a graph structure, which represents each phrase as nodes and the edges are similarity between nodes. Then, we proposed a new centrality computation method, which is based on the insight that the most important information typically occurs at the start or end of documents (document boundary) (Lin and Hovy, 1997; Teufel, 1997; Dong et al., 2021), to measure salience of local context based on the graph structure. Finally, we further combine the measure of global similarity and local salience for ranking. To evaluate the effectiveness of our method, we compare our method with recent state-of-the-art models on three public benchmarks (Inspec, DUC 2001, SemEval 2010). The results show that our model can outperform most models while generalizing better on input documents with different domains and length. It is deservedly mentioned that our models have a huge improvement on long scientific documents.

2 Methodology

The overall framework of our model is shown in Fig. 2. We follow the general process of unsuper-

vised keyphrase extraction. The main steps are as follows: (1) We tokenize the document and tag the document with part-of-speech (POS) tags. (2) We extract candidate phrases based on part-of-speech tags. We only keep noun phrases (NP) that consist of zero or more adjectives followed by one or multiple nouns (Wan and Xiao, 2008b). (3) We use a pre-trained language model to map the document text to low-dimension vector space and extract vector representation of candidate phrases and the whole document. (4) We score each candidate phrase with a rank algorithm which jointly models the global and local context. (5) We extract phrases with scores from the rank algorithm.

The main contribution of the whole process is the rank algorithm we proposed in step (4), which can be divided into three components: 1) phrase-document similarity for modeling global context; 2) boundary-aware centrality for modeling local context; 3) the combination of global and local information. We will introduce the details of these components in this section.

2.1 Document and Phrases Representations

Before introducing the rank algorithm, we first make clear step (1) - (3). We follow the common practice and use *StanfordCoreNLP Tools** to accomplish step (1) and (2). After previous universal steps, the document D was tokenized into tokens $\{t_1, t_2, \dots, t_N\}$ and candidate phrases $\{KP_0, KP_1, \dots, KP_n\}$ were extracted from document D . Different from previous works (Bennani-Smires et al., 2018) which use static vector to represent tokens in document, we employ BERT, which is a strong pre-trained language model, to obtain contextualized dynamic vector representations by Equ. (1).

$$\{H_1, H_2, \dots, H_N\} = \text{BERT}(\{t_1, t_2, \dots, t_N\}) \quad (1)$$

Where H_i is the vector representation of token t_i . Then, we obtain the vector representation H_{KP_i} of candidate phrases by computing the average of the phrase’s token vectors. The document vector representation is computed with max-pooling operation by Equ. (2).

$$H_D = \text{Maxpooling}(\{H_1, H_2, \dots, H_N\}) \quad (2)$$

Then, we can obtain document vector representation H_D which contains the global semantic information of the full document and a set of candidate

*<https://stanfordnlp.github.io/CoreNLP/>

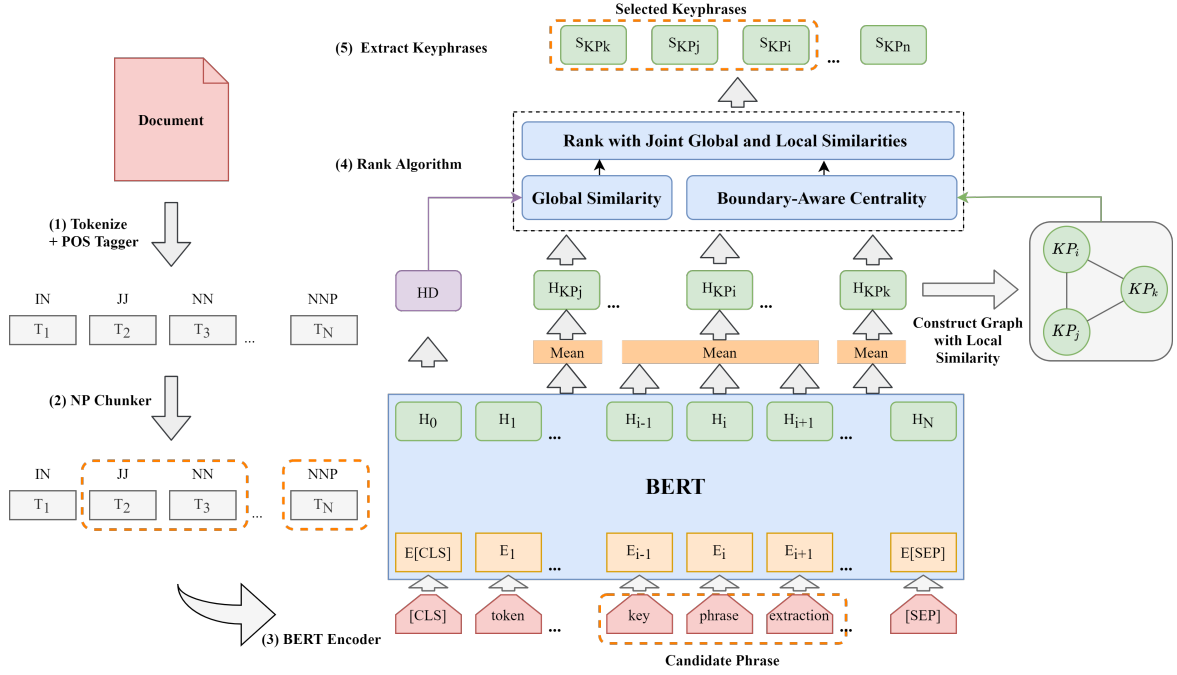


Figure 2: The framework of our unsupervised keyphrase extraction ranking model. (1) Get tokenized document and POS tags. (2) Extract noun phrases that consist of zero or more adjectives followed by one or multiple nouns. (3) Obtain embeddings of tokens in document with BERT. (4) Compute boundary-aware centrality and global relevance of each candidate phrases with global and local similarities. (5) Rank and extract keyphrases from candidate phrases with scores from the previous step.

phrase representations $V = \{H_{KP_i}\}_{i=1,\dots,n}$. Based on these representations, we will introduce the core rank algorithm of our model in the next section.

2.2 Proposed Rank Algorithm

2.2.1 Phrase-Document Similarity

We first introduce the computation of the phrase-document similarity for modeling global context. Specifically, we empirically employ Manhattan Distant (i.e. L1-distance) to compute similarity by Equ. (3).

$$\mathbf{R}(H_{KP_i}) = \frac{1}{\|H_D - H_{KP_i}\|_1} \quad (3)$$

Where $\|\cdot\|_1$ means Manhattan Distant and $R(H_{KP_i})$ represent the relevance between candidate phrase i and the whole document.

2.2.2 Traditional Degree Centrality

Graph-based ranking algorithms for keyphrase extraction represent a document as a graph $\mathcal{G} = (V, E)$, where $V = \{H_{KP_i}\}_{i=1,\dots,n}$ is the set of vector that represent nodes in graph (i.e. candidate phrases in document), and $E = \{e_{ij}\}$ is the set of edges that represent interactions between candidate phrases. In this paper, we simply employ the

degree of nodes as centrality to measure the importance of nodes. The degree centrality for candidate phrase i can be computed with Equ. (4).

$$C(H_{KP_i}) = \sum_{j=1}^n e_{ij} \quad (4)$$

Where $e_{ij} = H_{KP_i}^T \cdot H_{KP_j}$ is the dot-product similarity score for each pair (H_{KP_i}, H_{KP_j}) . We could also use other similarity measure methods (e.g. cosine similarity), but we empirically find that the simple dot-product performs better.

2.2.3 Boundary-Aware Centrality

Traditional centrality computation is based on the assumption that the contribution of the candidate phrase's importance in the document is not affected by the relative position of them, and the similarities of two graph nodes are symmetric. From human intuition, phrases that exist at the start or the end of a document should be more important than others. To implement this insight, we propose a new centrality computation method called boundary-aware centrality based on the assumption that important information typically occurs near boundaries (the start and end of documents) (Lin and Hovy, 1997; Teufel, 1997).

We reflect this assumption by employing a boundary function $d_b(i)$ over position of candidate phrases. This function d_b is formulated as Equ (5).

$$d_b(i) = \min(i, \alpha(n - i)) \quad (5)$$

Where n is the number of candidate phrases, and α is a hyper-parameter that controls relative importance of the start and end of a document. For node i and j , if $d_b(i) < d_b(j)$, then node i is closer to the boundary than node j . When calculating the centrality of node i , we need to reduce the contribution of node j to the centrality of node i . Based on this assumption and the boundary function $d_b(i)$, we can reconstruct the centrality computation of node i in the graph as Equ. (6).

$$\mathbf{C}(H_{\text{KP}_i}) = \sum_{d_b(i) < d_b(j)} e_{ij} + \lambda \sum_{d_b(i) \geq d_b(j)} e_{ij} \quad (6)$$

Where λ is used to reduce the influence of phrases which do not appear near the boundary to the centrality of node i .

Besides, we employ a threshold $\theta = \beta(\max(e_{ij}) - \min(e_{ij}))$ to filter the noise from nodes, which is far different from node i . We remove the influence of them to centrality by setting all $e_{ij} < \theta$ to zero. β is a hyper-parameter that controls the filter boundary. With the introduction of the noise filter strategy, we rewrite the Equ. (6) as Equ. (7).

$$\begin{aligned} \mathbf{C}(H_{\text{KP}_i}) = & \sum_{d_b(i) < d_b(j)} \max(e_{ij} - \theta, 0) \\ & + \lambda \sum_{d_b(i) \geq d_b(j)} \max(e_{ij} - \theta, 0) \end{aligned} \quad (7)$$

Where $\mathbf{C}(H_{\text{KP}_i})$ represents the local salience of candidate phrase i .

For most long documents or news articles, the author tends to write the key information at the beginning of the document. Florescu and Caragea (2017a) point out that the position-biased weight can greatly improve the performance for keyphrase extraction and they employ the sum of the position’s inverse of words in the document as the weight. For example, the word appearing at 2th, 5th and 10th, has a weight $p(w_i) = 1/2 + 1/5 + 1/10 = 0.8$. Our boundary-aware centrality has considered relative position information with boundary function. To prevent double counting, we follow a simpler position bias weight from (Sun

et al., 2020), which only considers where the candidate phrase first appears. The position bias weight is computed by $p(\text{KP}_i) = \frac{1}{p_1}$, where p_1 is the position of the candidate keyphrase’s first appearance. After that the softmax function is used to normalize the position bias weight as follow:

$$\hat{p}(\text{KP}_i) = \frac{\exp(p(\text{KP}_i))}{\sum_{k=1}^n \exp(p(\text{KP}_k))} \quad (8)$$

Then, boundary-aware centrality can be rewritten as Equ. (9).

$$\hat{\mathbf{C}}(H_{\text{KP}_i}) = \hat{p}(\text{KP}_i) \cdot \mathbf{C}(H_{\text{KP}_i}) \quad (9)$$

We finally employ $\hat{\mathbf{C}}(H_{\text{KP}_i})$ to measure the local salience of candidate phrase i .

2.2.4 Rank with Global and Local Information

To consider global and local level information at the same time, we simply combine the measure of global relevance $\mathbf{R}(H_{\text{KP}_i})$ and local salience $\hat{\mathbf{C}}(H_{\text{KP}_i})$ of candidate phrase together with multiplication to obtain the final score by Equ. (10).

$$\mathbf{S}(H_{\text{KP}_i}) = \mathbf{R}(H_{\text{KP}_i}) \cdot \hat{\mathbf{C}}(H_{\text{KP}_i}) \quad (10)$$

Finally, we rank candidate phrases with their final score $\mathbf{S}(H_{\text{KP}_i})$ and extract top-ranked k phrases as keyphrases of the document.

3 Experiments

3.1 Datasets and Evaluation Metrics

We evaluate our model on three public datasets: **Inspec**, **DUC2001** and **SemEval2010**. The **Inspec** dataset (Hulth, 2003) consists of 2,000 short documents from scientific journal abstracts. We follow previous works (Bennani-Smires et al., 2018; Sun et al., 2020) to use 500 test documents and the version of uncontrolled annotated keyphrases as ground truth. The **DUC2001** dataset (Wan and Xiao, 2008a) is a collection of 308 long length news articles with average 828.4 tokens. The **SemEval2010** dataset (Kim et al., 2010) contains ACM full length papers. In our experiments, we use the 100 test documents and the combined set of author- and reader- annotated keyphrases.

We follow the common practice and evaluate the performance of our models in terms of f-measure at the top N keyphrases (F1@N), and apply stemming to both extracted keyphrases and gold truth. Specifically, we report F1@5, F1@10 and F1@15 of each model on three datasets.

Models	DUC2001			Inspec			SemEval2010		
	F1@5	F1@10	F1@15	F1@5	F1@10	F1@15	F1@5	F1@10	F1@15
Statistical Models									
TF-IDF	9.21	10.63	11.06	11.28	13.88	13.83	2.81	3.48	3.91
YAKE	12.27	14.37	14.76	18.08	19.62	20.11	11.76	14.4	15.19
Graph-based Models									
TextRank	11.80	18.28	20.22	27.04	25.08	36.65	3.80	5.38	7.65
SingleRank	20.43	25.59	25.70	27.79	34.46	36.05	5.90	9.02	10.58
TopicRank	21.56	23.12	20.87	25.38	28.46	29.49	12.12	12.90	13.54
PositionRank	23.35	28.57	28.60	28.12	32.87	33.32	9.84	13.34	14.33
MultipartiteRank	23.20	25.00	25.24	25.96	29.57	30.85	12.13	13.79	14.92
Embedding-based Models									
EmbedRank d2v	24.02	28.12	28.82	31.51	37.94	37.96	3.02	5.08	7.23
EmbedRank s2v	27.16	31.85	31.52	29.88	37.09	38.40	5.40	8.91	10.06
SIFRank	24.27	27.43	27.86	29.11	38.80	39.59	-	-	-
SIFRank+	30.88	33.37	32.24	28.49	36.77	38.82	-	-	-
KeyGames	24.42	28.28	29.77	32.12	40.48	40.94	11.93	14.35	14.62
Proposed Model									
Our Model	28.62	35.52	36.29	32.61	40.17	41.09	13.02	19.35	21.72

Table 1: Comparison of our models with other baselines.

	DUC2001			Inspec			SemEval2010		
	F1@5	F1@10	F1@15	F1@5	F1@10	F1@15	F1@5	F1@10	F1@15
Our Model	28.62	35.52	36.29	32.49	40.04	41.05	12.26	19.22	21.42
- Global Similarity	26.22	33.85	34.61	30.75	38.49	40.52	12.00	18.93	21.29
- Local Similarity	17.45	18.64	19.03	22.51	27.58	30.36	8.81	10.7	11.45

Table 2: The results of ablation experiments on three datasets.

3.2 Comparison Models and Implementation Details

We compare our methods with three types of models to comprehensively prove the effectiveness of our models. Firstly, we compare with traditional statistical methods TF-IDF and YAKE (Campos et al., 2018). Secondly, We compare five strong graph-based ranking methods. TextRank (Mihalcea and Tarau, 2004) is the first attempt to convert text to graph with the co-occurrence of words and employ PageRank to rank phrases. SingleRank (Wan and Xiao, 2008a) improves the graph construction with a slide window. TopicRank (Bougouin et al., 2013) considers keyphrase extraction with topic distribution. PositionRank (Florescu and Caragea, 2017b) employs position information to weight the importance of phrases. MultipartiteRank (Boudin, 2018) splits the whole graph into sub-graph and ranks them with some graph theory. Finally, We compare three state-of-the-art embedding-based models. EmbedRank

(Bennani-Smires et al., 2018) first employs embedding of texts with Doc2Vec/Sent2Vec and measures the relevance of phrases and documents to select keyphrases. SIFRank (Sun et al., 2020) improves EmbedRank with contextualized embedding from a pre-trained language model. KeyGames (Saxena et al., 2020) creatively introduces game theoretic approach into automatic keyphrase extraction.

All the models use *Stanford CoreNLP Tools*[†] for tokenizing, part-of-speech tagging and noun phrase chunking. And regular expression $\{\langle NN.* | JJ \rangle * \langle NN.* \rangle\}$ is used to extract noun phrases as the candidate keyphrases. Our model’s hyperparameters for testing are chosen based on our results with the sampled 200 validation sets. The test results are chosen from the following hyperparameter settings: $\alpha \in \{0.5, 0.8, 1, 1.2, 1.5\}$, $\beta \in \{0.0, 0.1, 0.2, 0.3\}$ and $\lambda \in \{0.8, 0.9, 1.0\}$.

[†]<https://stanfordnlp.github.io/CoreNLP/>

3.3 Results

We report the results of our model in Tab. 1. We can observe that our model consistently outperforms most of the existing systems across the three datasets, each with different document length, covering two different domains. SIFRank and SIFRank+ have a remarkable performance on datasets with short input length due to document embedding of short documents can better represent the semantic information of full document and short document has fewer local information (e.g. aspects), which make embedding-based models perform well. We can further see that models with global similarity (i.e. EmbedRank and SIFRank) all outperform graph-based models on short length documents (i.e. DUC2001 and Inspec).

Compared with other works, our model and KeyGames, which is based on game theory, are more generalized and can tackle short and long input documents well. The advantages of our models are very obvious on the long scientific document dataset SemEval2010. This mainly benefits from the boundary-aware centrality for modeling local context.

3.4 Discussion

3.4.1 Ablation Study

We evaluate the contribution of the global and local component of our model with ablation study and the results can be seen in Tab. 2. From the results, we can find that the modeling of local context is more important than the modeling of global context. When we remove local information from our model, our model goes back to an embedding-based model. The performance on SemEval2010 is not sensitive to the removal of relevance-aware weighting. We guess that embedding of long documents may contain multi-aspects information which influences the measure of similarity between the phrase and the whole document, which leads to the influence of global information being limited. Overall, we can prove that jointly modeling global and local context is crucial for unsupervised keyphrase extraction and the revisit of degree centrality is effective for modeling local context and meaningful for future work.

3.4.2 Impact of Hyper-Parameters

In this section, we first analyze the impact of hyper-parameter β and then report the best setting on each dataset. We employ three hyper-parameters in our




Figure 3: Performance with different β .

models, β is used to filter noise and we can see the impact of β from Fig. 3. $\beta = 0.2$ is a proper configure for three datasets. α is used to control the importance of the start or end of a document. $\alpha < 1$ means the start of the document is more vital and $\alpha > 1$ means the end of the document is more vital.

The best settings are $\alpha = 0.8, \beta = 0.2, \lambda = 0.9$ on DUC2001, $\alpha = 0.5, \beta = 0.2, \lambda = 0.9$ on Inspec and $\alpha = 1.5, \beta = 0.2, \lambda = 0.8$ on SemEval2010. From these settings, we can get the following three conclusions, which is conforming to the characteristics of these datasets. For DUC2001 and Inspec, most vital information occurs at the start of the document due to the fact that DUC2001 is from news articles, and Inspec is from the abstract. For SemEval2010, the setting of α is contrary to previous datasets due to SemEval2010 is long scientific documents and much key information occurring at the end of the document (section conclusion). The settings of λ on three datasets show that long documents need to reduce more influence from contexts not near the boundary, which is intuitive.

3.4.3 Impact of Different Similarity Measure Methods

Our model employs Manhattan Distance to measure the similarity between phrases and the whole document. We also attempt to employ different measure methods. The results of different similarity measure methods are shown in Tab. 3, and we can see that the advantage of Manhattan Distance is obvious. We also can see that cosine similarity performs badly and is not suitable for our models.

Similarity Measure	DUC2001			Inspec			SemEval2010		
	F1@5	F1@10	F1@15	F1@5	F1@10	F1@15	F1@5	F1@10	F1@15
Euclidean Distance	23.31	28.04	30.39	28.5	37.01	29.25	10.99	16.37	18.41
Cosine similarity	15.01	17.96	19.44	23.67	30.26	33.35	9.70	12.22	13.29
Manhattan Distance	28.62	35.52	36.29	32.49	40.04	41.05	12.26	19.22	21.42

Table 3: The results of different measure methods for similarity between candidate phrase and the whole document.

Article	<p><u>Ben Johnson</u> should be stripped of his <u>world record</u> in the <u>100-meter dash</u> and it should be awarded to Carl Lewis, <u>many U.S. track</u> and field officials and athletes said Wednesday in the aftermath of Johnson admitting to his <u>steroid use</u>. Johnson testified this week at a <u>Canadian inquiry</u> that his seven-year involvement with <u>illegal performance-enhancing drugs</u> included injections before the 1987 <u>World Championships</u> in Rome, where he set the <u>existing world record</u> of 9.83 seconds. "I would have to see the evidence, but if he was on drugs at the time of the <u>World Championships</u>, my thought would be to remove his record," said Ollan Cassell, executive director of The <u>Athletics Congress</u>, the national governing board for the <u>sport</u>. Cassell is also a vice president of the <u>International Amateur Athletic Federation</u>, the world governing organization. The possibility that Johnson could lose his record was raised when the IAAF asked for "an urgent legal opinion" to see what it could do about records and championships of confessed drug users. The issue would be decided at the <u>IAAF Congress</u>, scheduled for ... There was much discussion among athletes competing this week in the <u>TAC Championships</u> about whether Johnson should be allowed back into the <u>sport</u> after he completes his two-year suspension. Joyner added. "He killed a lot of fans ... a lot of young kids ... and he almost killed our <u>sport</u> ... he put a dark shadow over it ... because of who he was," Joyner said ... Johnson was tested after his world-record run in Rome, but the results were negative.</p>
Gold Truth	<p><u>Ben Johnson</u>; <u>world record</u>; <u>100-meter dash</u>; <u>Carl Lewis</u>; <u>steroid use</u>; <u>Canadian inquiry</u>; <u>illegal performance-enhancing drugs</u>; <u>World Championships</u>;</p>
Our Model	<p>Top5: <u>world record</u>; <u>Ben Johnson</u>; <u>existing world record</u>; <u>World Championships</u>; International Amateur Athletic Federation; Top10: <u>100-meter dash</u>; <u>many U.S. track</u>; <u>legal time</u>; <u>world mark</u>; <u>Carl Lewis</u>; Top15: <u>sport</u>; <u>Athletics Congress</u>; <u>IAAF Congress</u>; <u>illegal performance-enhancing drugs</u>; <u>TAC Championships</u>;</p>

Figure 4: An example from DUC2001. The correct keyphrases are underlined. Red text means the extracted gold truth and blue text means extracted phrases by our model.

3.4.4 Case Study

In this section, we show an example from DUC2001 in Fig. 4. DUC2001 is a dataset from news articles. The correct keyphrases are underlined. Red text means the extracted gold truth and blue text means extracted phrases by our model.

We can see that all keyphrases occur at the start of the document. Our model extracted many correct phrases which are the same as gold truth and extracted the phrase "existing word record" which is semantically same with "word record" in gold truth. It is worth mentioning that our model focuses on the boundary of the document and most extracted phrases were located at the start of the document, which is controlled by our setting of α . This proves the effectiveness of our boundary-aware centrality. From the figure, we also can find that wrong phrases are highly relevant to topics of this document, which is influenced by our phrase-document relevance weighting. This example shows that the joint modeling of global and local context can improve the performance of keyphrase extraction and our model really captures local and global informa-

tion.

4 Related Work

4.1 Pre-trained Language Model

Pre-trained language model is the kind of model that is trained on large-scale unlabeled corpus to learn prior knowledge and then fine-tuned on downstream tasks. The pre-trained language model without fine-tuning also can provide high quality embedding of natural texts for unsupervised tasks. Different from static word embedding, such as Word2Vec (Mikolov et al., 2013), GloVe (Pennington et al., 2014), and FastText (Joulin et al., 2017). Pre-trained language models can encode words or sentences with context dynamically and solve the OOV problem. In addition, pre-trained language models can provide document-level or sentence-level embedding which contains more semantic information than Sen2Vec (Pagliardini et al., 2018) or Doc2Vec (Le and Mikolov, 2014).

ELMo (Peters et al., 2018) employs Bi-LSTM structure and concatenate forward and backward information to capture bidirectional information.

BERT (Devlin et al., 2019) is a bidirectional transformer structure pre-trained language model. Compared with the concatenation of bidirectional information, BERT can capture better context information. There are also a lot of other pre-trained language models such as RoBERTa (Liu et al., 2019), XLNET (Yang et al., 2020), etc. In this paper, we choose BERT, the most used, to obtain vector representation of documents and phrases by merging the embedding of tokens.

4.2 Unsupervised Keyphrase Extraction

Unsupervised keyphrase extraction can be divided into four main types: statistics-based models, graph-based models, topic-based models, and embedding-based models. Statistics-based models (Campos et al., 2018) mainly analyze an article’s probability features such as word frequency feature, position feature, linguistic features, etc. Topic-based models (Jardine and Teufel, 2014; Liu et al., 2009) focus on how to mine keyphrases by making use of the probability distribution of articles.

Graph-based models are the most proposed and popular used in early works which convert the document into a graph. Inspired by (Page et al., 1999), (Mihalcea, 2004) proposed TextRank to convert keyphrase extraction task into the rank of nodes in graph. After this, various works focused on the expansion of TextRank. (Wan and Xiao, 2008a) proposed SingleRank, which employs co-occurrences of tokens as edge weights. (Bougouin et al., 2013) proposed TopicRank, which assigns a significance score to each topic by candidate keyphrase clustering. MultipartiteRank (Boudin, 2018) encodes topical information within a multipartite graph structure. Recently, (Wang, 2015) proposed WordAttractionRank, which added distance between word embeddings into SingleRank, and (Florescu and Caragea, 2017b) use node position weights, favoring words appearing earlier in the text. This position bias weighting strategy is very useful in news articles and long documents.

Embedding-based models benefit from the development of representation learning, which maps natural language into low-dimension vector representation. Therefore, in recent years, embedding-based keyphrase extraction (Wang et al., 2016; Bennani-Smires et al., 2018; Papagiannopoulou and Tsoumakas, 2018; Sun et al., 2020) has achieved good performance. (Bennani-Smires et al., 2018) proposed EmbedRank, which ranks

phrases by measuring the similarity between phrase embedding and document embedding. (Sun et al., 2020) proposed SIFRank, which improves the static embedding from EmbedRank with a pre-trained language model.

Embedding-based models just measured the similarity between document and candidate phrases and ignored the local information. To jointly model global and local context (Zheng and Lapata, 2019; Liang et al., 2021), in this paper, we revisit degree centrality, which can model local context, and convert it into boundary-aware centrality. Then, we combine global similarity and boundary-aware centrality for local salient information to rank and extract phrases.

5 Conclusion and Future Work

In this paper, we point out that embedding-based models ignore the local information and propose a novel model which jointly models global and local context. Our model revisited degree centrality and modified it with boundary function for modeling local context. We combine global similarity with our proposed boundary-aware centrality to extract keyphrases. Experiments on 3 public benchmarks demonstrate that our model can effectively capture global and local information and achieve remarkable results. In the future work, we will focus on how to introduce our boundary-aware mechanism into supervised end2end keyphrase extraction/generation models.

Acknowledgments

We thank the three anonymous reviewers for their careful reading of our paper and their many insightful comments and suggestions. This work was supported in part by the National Natural Science Foundation of China (Grant Nos.U1636211, 61672081,61370126), the 2020 Tencent Wechat Rhino-Bird Focused Research Program, and the Fund of the State Key Laboratory of Software Development Environment (Grant No. SKLSDE-2021ZX-18).

References

Kamil Bennani-Smires, Claudiu Musat, Andreea Hossmann, Michael Baeriswyl, and Martin Jaggi. 2018. [Simple unsupervised keyphrase extraction using sentence embeddings](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learn-*

- ing, pages 221–229, Brussels, Belgium. Association for Computational Linguistics.
- Florian Boudin. 2018. [Unsupervised keyphrase extraction with multipartite graphs](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 667–672, New Orleans, Louisiana. Association for Computational Linguistics.
- Adrien Bougouin, Florian Boudin, and Béatrice Daille. 2013. [TopicRank: Graph-based topic ranking for keyphrase extraction](#). In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 543–551, Nagoya, Japan. Asian Federation of Natural Language Processing.
- Ricardo Campos, Vítor Mangaravite, Arian Pasquali, Alípio Mário Jorge, Célia Nunes, and Adam Jatowt. 2018. [Yake! collection-independent automatic keyword extractor](#). In *Advances in Information Retrieval*, pages 806–810, Cham. Springer International Publishing.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).
- Yue Dong, Andrei Mircea, and Jackie C. K. Cheung. 2021. [Discourse-aware unsupervised summarization of long scientific documents](#).
- Corina Florescu and Cornelia Caragea. 2017a. [A position-biased pagerank algorithm for keyphrase extraction](#).
- Corina Florescu and Cornelia Caragea. 2017b. [PositionRank: An unsupervised approach to keyphrase extraction from scholarly documents](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1105–1115, Vancouver, Canada. Association for Computational Linguistics.
- Kazi Saidul Hasan and Vincent Ng. 2014. [Automatic keyphrase extraction: A survey of the state of the art](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1262–1273, Baltimore, Maryland. Association for Computational Linguistics.
- Anette Hulth. 2003. [Improved automatic keyword extraction given more linguistic knowledge](#). In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*, pages 216–223.
- James Jardine and Simone Teufel. 2014. [Topical PageRank: A model of scientific expertise for bibliographic search](#). In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 501–510, Gothenburg, Sweden. Association for Computational Linguistics.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. [Bag of tricks for efficient text classification](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431, Valencia, Spain. Association for Computational Linguistics.
- Su Nam Kim, Olena Medelyan, Min-Yen Kan, and Timothy Baldwin. 2010. [SemEval-2010 task 5 : Automatic keyphrase extraction from scientific articles](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 21–26, Uppsala, Sweden. Association for Computational Linguistics.
- Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32, ICML'14*, page II–1188–II–1196. JMLR.org.
- Xinnian Liang, Shuangzhi Wu, Mu Li, and Zhoujun Li. 2021. [Improving unsupervised extractive summarization with facet-aware modeling](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1685–1697, Online. Association for Computational Linguistics.
- Chin-Yew Lin and Eduard Hovy. 1997. [Identifying topics by position](#). In *Fifth Conference on Applied Natural Language Processing*, pages 283–290, Washington, DC, USA. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Zhiyuan Liu, Peng Li, Yabin Zheng, and Maosong Sun. 2009. [Clustering to find exemplar terms for keyphrase extraction](#). In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 257–266, Singapore. Association for Computational Linguistics.
- Rada Mihalcea. 2004. [Graph-based ranking algorithms for sentence extraction, applied to text summarization](#). In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 170–173, Barcelona, Spain. Association for Computational Linguistics.
- Rada Mihalcea and Paul Tarau. 2004. [TextRank: Bringing order into text](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 404–411, Barcelona, Spain. Association for Computational Linguistics.
- Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#). In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.

- Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab.
- Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. 2018. Unsupervised Learning of Sentence Embeddings using Compositional n-Gram Features. In *NAACL 2018 - Conference of the North American Chapter of the Association for Computational Linguistics*.
- Eirini Papagiannopoulou and Grigorios Tsoumakas. 2018. Local word vectors guiding keyphrase extraction.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Arnav Saxena, Mudit Mangal, and Goonjan Jain. 2020. KeyGames: A game theoretic approach to automatic keyphrase extraction. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2037–2048, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Y. Sun, H. Qiu, Y. Zheng, Z. Wang, and C. Zhang. 2020. Sifrank: A new baseline for unsupervised keyphrase extraction based on pre-trained language model. *IEEE Access*, 8:10896–10906.
- Simone Teufel. 1997. Sentence extraction as a classification task. In *Intelligent Scalable Text Summarization*.
- Xiaojun Wan and Jianguo Xiao. 2008a. Single document keyphrase extraction using neighborhood knowledge. In *Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 2, AAAI'08*, page 855–860. AAAI Press.
- Xiaojun Wan and Jianguo Xiao. 2008b. Single document keyphrase extraction using neighborhood knowledge. In *Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 2, AAAI'08*, page 855–860. AAAI Press.
- R. Wang. 2015. Corpus-independent generic keyphrase extraction using word embedding vectors.
- Yunli Wang, Yong Jin, Xiaodan Zhu, and Cyril Goutte. 2016. Extracting discriminative keyphrases with learned semantic hierarchies. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 932–942, Osaka, Japan. The COLING 2016 Organizing Committee.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2020. Xlnet: Generalized autoregressive pretraining for language understanding.
- Hao Zheng and Mirella Lapata. 2019. Sentence centrality revisited for unsupervised summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6236–6247, Florence, Italy. Association for Computational Linguistics.