

# Importance Estimation from Multiple Perspectives for Keyphrase Extraction

Mingyang Song, Liping Jing\* and Lin Xiao

Beijing Key Lab of Traffic Data Analysis and Mining

Beijing Jiaotong University, China

{mingyang.song, lpjing, 17112079}@bjtu.edu.cn

## Abstract

Keyphrase extraction is a fundamental task in Natural Language Processing, which usually contains two main parts: candidate keyphrase extraction and keyphrase importance estimation. From the view of human understanding documents, we typically measure the importance of phrase according to its syntactic accuracy, information saliency, and concept consistency simultaneously. However, most existing keyphrase extraction approaches only focus on the part of them, which leads to biased results. In this paper, we propose a new approach to estimate the importance of keyphrase from multiple perspectives (called as *KIEMP*) and further improve the performance of keyphrase extraction. Specifically, *KIEMP* estimates the importance of phrase with three modules: a chunking module to measure its syntactic accuracy, a ranking module to check its information saliency, and a matching module to judge the concept (i.e., topic) consistency between phrase and the whole document. These three modules are seamlessly jointed together via an end-to-end multi-task learning model, which is helpful for three parts to enhance each other and balance the effects of three perspectives. Experimental results on six benchmark datasets show that *KIEMP* outperforms the existing state-of-the-art keyphrase extraction approaches in most cases.

## 1 Introduction

Keyphrase Extraction (KE) aims to select a set of reliable phrases (e.g., “harmonic balance method”, “groebner base”, “error bound”, “algebraic representation”, and “singular point” in Table 1) with salient information and central topics from a given document, which is a fundamental task in natural language processing. Most classic keyphrase extraction methods typically include two mainly components: *candidate keyphrase extraction* and

*Input Document:*

*harmonic balance (hb) method* is well known principle for analyzing periodic oscillations on nonlinear networks and systems. because the hb method has a truncation error, approximated solutions have been guaranteed by *error bounds*. however, its numerical computation is very time consuming compared with solving the hb equation. this paper proposes an *algebraic representation* of the *error bound* using *groebner base*. the *algebraic representation* enables to decrease the computational cost of the *error bound* considerably. moreover, using *singular points* of the *algebraic representation*, we can obtain accurate break points of the *error bound* by collisions.

*Output / Target Keyphrases:*

*harmonic balance method; groebner base; error bound; algebraic representation; singular point; quadratic approximation*

Table 1: Sample input document with output / target keyphrases in *KP20k* testing set. Specially, keyphrases typically can be categorized into two types: *present keyphrase* that appears in a given document and *absent keyphrase* which does not appear in a given document.

*keyphrase importance estimation* (Medelyan et al., 2009; Liu et al., 2010; Hasan and Ng, 2014).

As shown in Table 1, each keyphrase usually consists of more than one words (Meng et al., 2017). To extract the candidate keyphrases from the given document which is typically characterized via word-level representation, researchers leverage some heuristics (Wan and Xiao, 2008; Liu et al., 2009a,b; Nguyen and Phan, 2009; Grineva et al., 2009; Medelyan et al., 2009) to identify the candidate keyphrases. For example, the word embeddings are composed to n-grams by Convolution Neural Network (CNN) (Xiong et al., 2019; Sun et al., 2020; Wang et al., 2020).

Usually, the candidate set contains much more keyphrases than the ground truth keyphrase set. Therefore, it is critical to select the important keyphrase from the candidate set. In other words, *keyphrase importance estimation* commonly is one of the essential components in many keyphrase extraction models. Since the keyphrase extraction

\*Corresponding author.

task concerns “the automatic selection of important and topical phrases from the body of a document” (Turney, 2000). In other words, its goal is to estimate the importance of the candidate keyphrases to determine which one should be extracted. Recent approaches (Sun et al., 2020; Wang et al., 2020) recast the keyphrase extraction as a classification problem, which extracts keyphrases by a binary classifier. However, a binary classifier classifies each candidate keyphrase independently, and consequently, it does not allow us to determine which candidates are better than the others (Hulth, 2004). Therefore, some methods (Jiang et al., 2009; Xiong et al., 2019; Wang et al., 2020; Sun et al., 2020) propose a ranking model to extract keyphrases, where the goal is to learn a phrase ranker to compare the saliency of two candidate phrases. Furthermore, many previous studies (Liu et al., 2010; Wang et al., 2019; Liu et al., 2009b) extract keyphrases with the main topics discussed in the source document. For example, Liu et al. (2010) proposes to build a topical PageRank approach to measure the importance of words concerning different topics.

However, most existing keyphrase extraction methods estimate the importance of keyphrases on at most two perspectives, leading to biased extraction. Therefore, to improve the performance of keyphrase extraction, the importance of the candidate keyphrases requires to be estimated sufficiently from multiple perspectives. Motivated by the phenomenon mentioned above, we propose a new importance estimation from multiple perspectives simultaneously for the keyphrase extraction task. Concretely, it estimates the importance from three perspectives with three modules (syntactic accuracy, information saliency, and concept consistency) with three modules. A chunking module, as a binary classification layer, measures the syntactic accuracy of each candidate keyphrase. A ranking module checks the semantics saliency of each candidate phrase by a pairwise ranking approach, which introduces competition between the candidate keyphrases to extract more salient keyphrases. A matching module judges the concept relevance of each candidate phrase in the document via a metric learning framework. Furthermore, our model is trained jointly on the above three modules, balancing the effect of three perspectives. Experimental results on two benchmark data sets show that *KIEMP* outperforms the existing state-of-the-art keyphrase extraction approaches in most cases.

## 2 Related Work

A good keyphrase extraction system typically consists of two steps: (1) *candidate keyphrase extraction*, extracting a list of words / phrases that serve as the candidate keyphrases using some heuristics (Wan and Xiao, 2008; Nguyen and Phan, 2009; Medelyan et al., 2009; Grineva et al., 2009; Liu et al., 2009a,b); and (2) *keyphrase importance estimation*, determining which of these candidate phrases are keyphrases using different importance estimation approaches.

In the *candidate keyphrase extraction*, the heuristic rules usually are designed to avoid spurious phrases and keep the number of candidates to a minimum (Hasan and Ng, 2014). Generally, the heuristics mainly include (1) leverage a stop word list (Liu et al., 2009b), (2) allowing words with part-of-speech tags (Mihalcea and Tarau, 2004; Liu et al., 2009a), (3) composing words to n-grams to be the candidate keyphrases (Medelyan et al., 2009; Sun et al., 2020; Xiong et al., 2019; Wang et al., 2020). The above heuristics have proven effective with their high recall in extracting gold keyphrases from various sources. Motivated by the above methods, in this paper, we leverage CNNs to compose words to n-grams as the candidate keyphrases.

In the *keyphrase importance estimation*, the existing methods can be mainly divided into two categories: unsupervised and supervised. The unsupervised method usually are categorized into four groups, i.e., graph-based ranking (Mihalcea and Tarau, 2004), topic-based clustering (Liu et al., 2009b), simultaneous learning (Zha, 2002), and language modeling (Tomokiyo and Hurst, 2003). Early supervised approaches to keyphrase extraction recast this task as a binary classification problem (Witten et al., 1999; Turney, 2002, 2000; Jiang et al., 2009). Later, to determine which candidates are better than the others, many ranking approach is proposed to rank the saliency of two phrases (Jiang et al., 2009; Sun et al., 2020). This pairwise ranking approach, therefore, introduces competition between candidate keyphrases and has been achieved good performance. Both supervised and unsupervised methods construct features or models from different perspectives to measure the importance of candidate keyphrases to determine which keyphrases should be extracted. However, the approaches mentioned earlier consider at most two perspectives when measuring the importance of phrases, which leads to biased keyphrase extraction.

Different from the existing methods, the proposed *KIEMP* considers estimating the importance of the candidate keyphrases from multiple perspectives simultaneously.

### 3 Methodology

We formally define the problem of keyphrase extraction as follows. In this paper, *KIEMP* takes a document  $D = \{w_1, \dots, w_i, \dots, w_M\}$  and learns to extract a set of keyphrases  $K$  (each keyphrase may be composed of one or several word(s)) from their  $n$ -gram based representations under multiple perspectives.

This section describes the architecture of *KIEMP*, as shown in Figure 1. *KIEMP* mainly consists of two submodels: candidate keyphrase extraction and keyphrase importance estimation. The former first identifies and extracts the candidate keyphrases. Then the latter estimates the importance of keyphrases from three perspectives simultaneously with three modules to determine which one should be extracted.

#### 3.1 Contextualized Word Representation

Recently, pre-trained language models (Peters et al., 2018; Devlin et al., 2019; Liu et al., 2019) have emerged as a critical technology for achieving impressive gains in a wide variety of natural language tasks (Liu and Lapata, 2019). These models extend the idea of word embeddings by learning contextual representations from large-scale corpora using a language modeling objective. In this situation, Xiong et al. (2019) propose to represent each word by its ELMo (Peters et al., 2018) embedding and Sun et al. (2020) leverage variants of BERT (Devlin et al., 2019; Liu et al., 2019) to obtain contextualized word representations. Motivated by the above approaches, we represent each word by RoBERTa (Liu et al., 2019), which encodes  $D$  to a sequence of vector  $H = \{h_1, \dots, h_i, \dots, h_M\}$ :

$$H = \text{RoBERTa}\{w_1, \dots, w_i, \dots, w_M\}, \quad (1)$$

where  $h_i \in \mathbb{R}^d$  indicates the  $i$ -th contextualized word embedding of  $w_i$  from the last transformer layer in RoBERTa. Specifically, the [CLS] token of RoBERTa is used as the document representation.

#### 3.2 Candidate Keyphrase Extraction

In the keyphrase extraction task, keyphrase usually contains more than one word, as shown in Table 1. Therefore, it is necessary to identify the

candidate keyphrases via some strategies. Previous work (Medelyan et al., 2009; Sun et al., 2020; Wang et al., 2020; Xiong et al., 2019) allow  $n$ -grams that appear in the document to be the candidate keyphrases. Motivated by the previous approaches, we consider the language properties (Xiong et al., 2019) and compose the contextualized word representations to  $n$ -grams by CNNs (similar to Sun et al. (2020)). Specifically, the phrase representation of the  $i$ -th  $n$ -gram  $c_i^n$  is computed as:

$$h_i^n = \text{CNN}^n(h_{i:i+n}), \quad (2)$$

where  $h_i^n \in \mathbb{R}^d$  indicates the  $i$ -th  $n$ -gram representation. Concretely,  $n \in [1, N]$  is the length of  $n$ -grams, and  $N$  indicates the maximum length of allowed candidate  $n$ -grams. Specifically, each  $n$ -gram has its own set of convolution filters  $\text{CNN}^n$  with window size  $n$  and stride 1.

#### 3.3 Keyphrase Importance Estimation

In the keyphrase extraction models, *keyphrase importance estimation* commonly is one of the essential components. To improve the accuracy of keyphrase extraction, we estimate the importance of keyphrases from three perspectives simultaneously with three modules: chunking for syntactic accuracy, ranking for information saliency, and matching for concept consistency.

##### 3.3.1 Chunking for Syntactic Accuracy

Many studies (Turney, 2002; Witten et al., 1999; Turney, 2000) regard keyphrase extraction as a classification task, in which a model is trained to determine whether a candidate phrase is a keyphrase in a syntactic perspective. For example, Xiong et al. (2019); Sun et al. (2020) directly predict whether the  $n$ -gram is a keyphrase based on its corresponding representation. Motivated by these above methods, in this paper, the syntactic accuracy of phrase  $c_i^n$  is estimated by a chunking module:

$$I_1(c_i^n) = \text{softmax}(\mathbf{W}_1 h_i^n + b_1), \quad (3)$$

where  $\mathbf{W}_1$  and  $b_1$  indicate a trainable matrix and a bias. The softmax is taken over all possible  $n$ -grams at each position  $i$  and each length  $n$ . The whole model is trained using cross-entropy loss:

$$L_c = \text{CrossEntropy}(y_i^n, I_1(c_i^n)), \quad (4)$$

where  $y_i^n$  is the label of whether the phrase  $c_i^n$  is a keyphrase of the original document.

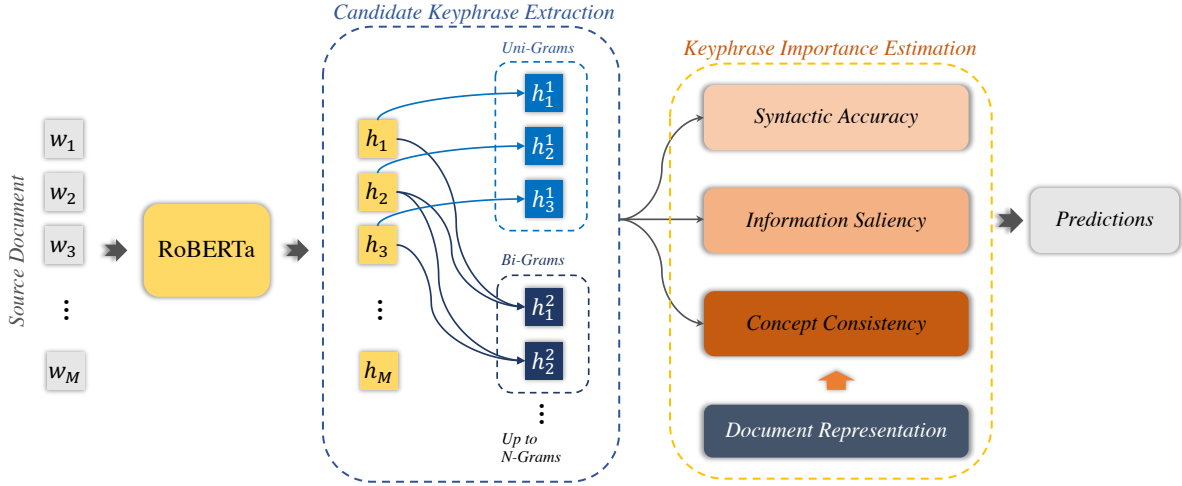


Figure 1: The *KIEMP* model architecture.

### 3.3.2 Ranking for Information Saliency

The binary classifier-based keyphrase extraction model classifies each candidate keyphrase independently, and consequently, it does not allow us to determine which candidates are better than the others (Hulth, 2004). However, the goal of keyphrase extraction is to identify the most salient phrases for a document (Hasan and Ng, 2014). Therefore, a ranking model is required to rank the saliency of the candidate keyphrases. We leverage a pairwise learning approach to rank the candidate keyphrases globally to compare the information saliency between all candidates. First, we put the candidate keyphrases in the document that are labeled as keyphrases, in the positive set  $\mathbf{P}^+$ , and the others to the negative set  $\mathbf{P}^-$ , to obtain the ranking labels. Then, the loss function is the standard hinge loss in the pairwise learning model:

$$L_r = \sum_{p^+, p^- \in K} \max(0, \delta_1 - I_2(p^+) + I_2(p^-)), \quad (5)$$

where  $I_2(\cdot)$  represents the estimation of information saliency and  $\delta_1$  indicates the margin. It enforces *KIEMP* to rank the candidate keyphrases  $p^+$  ahead of  $p^-$  within the same document. Specifically, the information saliency of the  $i$ -th  $n$ -gram representation  $c_i^n$  can be computed as follows:

$$I_2(c_i^n) = \mathbf{W}_2 h_i^n + b_2, \quad (6)$$

where  $\mathbf{W}_2$  is a trainable matrix, and  $b_2$  is a bias. Through the pairwise learning model, we can rank the information saliency of all candidates and extract the keyphrases with more salient information sufficiently.

### 3.3.3 Matching for Concept Consistency

As phrases are used to express various meanings corresponding to different concepts (i.e., topics), a phrase will play different important roles in different concepts of the document (Liu et al., 2010). A matching module is proposed via a metric learning framework to estimate the concept consistency between the candidate keyphrases and their corresponding document. We first apply variation autoencoder (Rezende et al., 2014) on the documents  $\mathbf{D}$  and the candidate keyphrases  $\mathbf{K}$  to obtain their concepts. Each document  $D$  is encoded via a latent variable  $z \in \mathbb{R}^c$  which is assumed to be sampled from a standard Gaussian prior, i.e.,  $z \sim p(z) = \mathcal{N}(0, I_d)$ . Such variable has ability to determine the latent concepts hidden in the documents and will be useful to extract keyphrase (Wang et al., 2019). During the encoding process,  $z$  can be sampled via a re-parameterization trick for Gaussian distribution, i.e.,  $z \sim q(z|D) = \mathcal{N}(\mu, \sigma)$ . Specifically, we sample an auxiliary noise variable  $\varepsilon \sim \mathcal{N}(0, I)$  and re-parameterization  $z = \mu + \sigma \odot \varepsilon$ , where  $\odot$  denotes the element-wise multiplication. The mean vector  $\mu \in \mathbb{R}^c$  and variance vector  $\sigma \in \mathbb{R}^c$  will be inferred by a two-layer network with ReLU-activated function, i.e.,  $\mu = \mu_\phi(D)$  and  $\sigma = \sigma_\phi(D)$  where  $\phi$  is the parameter set. During the decoding process, the document can be reconstructed by a multi-layer network ( $f_k$ ) with Tanh-activated function, i.e.,  $\tilde{D} = f_k(z)$ . Furthermore, the candidate keyphrases are processed in the same way as the documents.

Once having the latent concept representation

of the document  $z$  and the phrase  $z_i^n$ , the concept consistency can be estimated as follows,

$$I_3(c_i^n, D) = z_i^n \mathbf{W}_3 z. \quad (7)$$

Here,  $\mathbf{W}_3$  is a learnable mapping matrix. The loss function is the triplet loss in the metric learning framework calculated as follows:

$$L_m = \sum_{p^+, p^- \in K} \max(0, I_3(p^-, D) - I_3(p^+, D) + \delta_2), \quad (8)$$

where  $\delta_2$  represents the margin. It enforces *KIEMP* to match and rank the concept consistency of keyphrases  $p^+$  ahead of the non-keyphrases  $p^-$  within their corresponding document  $D$ .

Furthermore, to simultaneously minimize the reconstruction loss and penalize the discrepancy between a prior distribution and posterior distribution about the latent variable  $z$ , the VAE process can be implemented by optimizing the following objective function for the documents  $L_d$  and the candidate keyphrases  $L_k$ :

$$L_d = -\mathbb{E}_{q(\mathbf{z}|\mathbf{D})} [p(\mathbf{D}|\mathbf{z})] + D_{KL}(p(\mathbf{z})||q(\mathbf{z}|\mathbf{D})), \quad (9)$$

$$L_k = -\mathbb{E}_{q(\mathbf{z}|\mathbf{K})} [p(\mathbf{K}|\mathbf{z})] + D_{KL}(p(\mathbf{z})||q(\mathbf{z}|\mathbf{K})), \quad (10)$$

where  $D_{KL}$  indicates the Kullback-Leibler divergence between two distributions. And the final loss of this module is calculated as follows:

$$L_t = L_m + \lambda L_d + (1 - \lambda)L_k, \quad (11)$$

where  $\lambda \in (0, 1)$  is the balance factor. Through concept consistency matching, we expect to align keyphrases with high-level concepts (i.e., topics or structures) in the document to assist the model in extracting keyphrases with more important concepts.

### 3.4 Model Training and Inference

Multi-task learning has played an essential role in various fields (Srna et al., 2018), and has been widely used in the natural language processing tasks (Sun et al., 2020; Mu et al., 2020) recently. Therefore, our framework allows end-to-end learning of syntactic chunking, saliency ranking, and concept matching in this paper. Then, we define the training objective of the entire framework with the linear combination of  $L_c$ ,  $L_r$ , and  $L_t$ :

$$L = \epsilon_1 L_c + \epsilon_2 L_r + \epsilon_3 L_t, \quad (12)$$

where the hyper-parameters  $\epsilon_1$ ,  $\epsilon_2$ , and  $\epsilon_3$  balance the effects of the importance estimation from three perspectives. Specifically,  $\epsilon_1 + \epsilon_2 + \epsilon_3 = 1$ .

In this paper, *KIEMP* aims to extract keyphrases according to their saliency. It contains three modules syntactic accuracy chunking, information saliency ranking, and concept consistency matching. Chunking and matching are used to enforce the ranking module to rank the proper candidate keyphrases ahead. Therefore, only the ranking module is used in the inference process (test-phase).

Dataset	Document Len.	# Keyphrase	Keyphrase Len.
	Average	Average	Average
<i>OpenKP</i>	900.4	1.8	2.0
<i>KP20k</i>	179.8	5.3	2.0
<i>Inspec</i>	128.7	9.8	2.5
<i>Krapivin</i>	182.6	5.8	2.2
<i>Nus</i>	219.1	11.7	2.2
<i>SemEval</i>	234.8	14.7	2.4

Table 2: Statistics of six benchmark datasets. *Document Len.* and *Keyphrase Len.* represent the number of words in the document and keyphrase respectively.

## 4 Experimental Settings

### 4.1 Datasets

Six benchmark datasets are mainly used in our experiments, *OpenKP* (Xiong et al., 2019), *KP20k* (Meng et al., 2017), *Inspec* (Hulth, 2003), *Krapivin* (Krapivin and Marchese, 2009), *Nus* (Nguyen and Kan, 2007) and *SemEval* (Kim et al., 2010). Table 2 summarizes the statistics of each testing sets.

*OpenKP* consists of around 150K documents sampled from the index of the Bing search engine. In *OpenKP*, we follow the official split of training (134K documents), development (6.6K documents), and testing (6.6K documents) sets. The keyphrases for each document in *OpenKP* were labeled by expert annotators, with each document assigned 1-3 keyphrases. As a requirement, all the keyphrases appeared in the original document (Xiong et al., 2019).

*KP20k* contains a large number of high-quality scientific metadata in the computer science domain from various online digital libraries (Meng et al., 2017). We follow the official setting of this dataset and split the dataset into training (528K

Hyper-parameter	Dimension or Value
$\lambda$	0.5
$\epsilon_1, \epsilon_2, \epsilon_3$	1/3
$\delta_1, \delta_2$	1.0
Optimizer	AdamW
Learning Rate	$1 \times 10^{-5}$
Batch Size	32
Warm-Up Proportion	10%
RoBERTa Embedding ( $\mathbb{R}^d$ )	768
Concept Dimension ( $\mathbb{R}^c$ )	64
Max Sequence Length	512
Maximum Phrase Length ( $N$ )	5

Table 3: Parameters used for training *KIEMP*.

documents), validation (20K documents), and testing (20K documents) sets. From the training set of *KP20k*, we remove all articles that are duplicated in themselves, either in the *KP20k* validation and testing set. After the cleanup, the *KP20k* dataset contains 504K training samples, 20K validation samples, and 20K testing samples.

To verify the robustness of *KIEMP*, we also test the model trained with *KP20k* dataset on four widely-adopted keyphrase extraction data sets including *Inspec*, *Krapivin*, *Nus*, and *SemEval*.

In this paper, we focus on keyphrase extraction. Therefore, only the keyphrases that appear in the documents are used for training and evaluation.

## 4.2 Baselines

This paper focuses on the comparisons with the state-of-the-art baselines and chooses the following keyphrase extraction models as our baselines.

**TextRank** An unsupervised algorithm based on weighted-graphs proposed by [Mihalcea and Tarau \(2004\)](#). Given a word graph built on co-occurrences, it calculates the importance of candidate words with PageRank. The importance of a candidate keyphrase is then estimated as the sum of the scores of the constituent words.

**TFIDF** ([Jones, 2004](#)) is computed based on candidate frequency in the given text and inverse document frequency

**CopyRNN** ([Meng et al., 2017](#)) which uses the attention mechanism as the copy mechanism to extract keyphrases from the given document.

**BLING-KPE** ([Xiong et al., 2019](#)) first concatenates the pre-trained language model (ELMo ([Peters et al., 2018](#))) as word embeddings, visual as

well as positional features, and then uses a CNN network to obtain n-gram phrase embeddings for binary classification.

**JointKPE** ([Sun et al., 2020](#)) jointly learns a chunking model (*ChunkKPE*) and a ranking model (*RankKPE*) for keyphrase extraction.

**SMART-KPE+R2J** ([Wang et al., 2020](#)) presents a multi-modal method to the keyphrase extraction task, which leverages lexical and visual features to enable strategy induction as well as meta-level features to aid in strategy selection.

**DivGraphPointer** ([Sun et al., 2019](#)) combines the advantages of traditional graph-based ranking methods and recent neural network-based approaches. Furthermore, they also propose a diversified point network to generate a set of diverse keyphrases out of the word graph in the decoding process.

**Div-DGCN** ([Zhang et al., 2020](#)) proposes to adopt the Dynamic Graph Convolutional Networks (DGCN) to acquire informative latent document representation and better model the compositionality of the target keyphrases set.

**SKE-Large-CLS** ([Mu et al., 2020](#)) obtains span-based representation for each keyphrase and further learns to capture the similarity between keyphrases in the source document to get better keyphrase predictions.

In this paper, for ease of introduction, all the baselines are divided according to the following three perspectives, syntax, saliency, and combining syntax and saliency. Among them, *BLING-KPE*, *CopyRNN*, *ChunkKPE* belong to the former, *TFIDF*, *TextRank*, as well as *RankKPE* belong to the second, and *DivGraphPointer*, *Div-DGCN*, *SKE-Large-CLS*, *SMART-KPE+R2J*, as well as *JointKPE* belongs to the last.

## 4.3 Evaluation Metrics

For the keyphrase extraction task, the performance of keyphrase model is typically evaluated by comparing the top  $k$  predicted keyphrases with the target keyphrases (ground-truth labels). The evaluation cutoff  $k$  can be a fixed number (e.g.,  $F_1@5$  compares the top-5 keyphrases predicted by the model with the ground-truth to compute an  $F_1$  score). Following the previous work ([Meng et al., 2017](#); [Sun et al., 2019](#)), we adopt macro-averaged recall and F-measure ( $F_1$ ) as evaluation metrics, and  $k$  is set to be 1, 3, 5, and 10. In the evaluation, we apply Porter Stemmer ([Porter, 2006](#)) to both

<i>Model</i>	<i>OpenKP</i>						<i>KP20k</i>	
	<i>R@1</i>	<i>R@3</i>	<i>R@5</i>	<i>F<sub>1</sub>@1</i>	<i>F<sub>1</sub>@3</i>	<i>F<sub>1</sub>@5</i>	<i>F<sub>1</sub>@5</i>	<i>F<sub>1</sub>@10</i>
<i>Unsupervised Methods</i>								
<i>TFIDF (Jones, 2004)</i>	<u>0.150</u>	<u>0.284</u>	<u>0.347</u>	0.196*	0.223*	0.196*	0.105	0.130
<i>TextRank (Mihalcea and Tarau, 2004)</i>	0.041	0.098	0.142	0.054*	0.076*	0.079*	<u>0.180</u>	<u>0.150</u>
<i>Supervised Methods with Additional Features</i>								
<i>BLING-KPE (Xiong et al., 2019)</i>	0.220	0.390	0.481	0.285*	0.303*	0.270*	-	-
<i>SMART-KPE+R2J (Wang et al., 2020)</i>	<u>0.307</u>	<u>0.532</u>	<u>0.625</u>	<u>0.381</u>	<u>0.405</u>	<u>0.347</u>	-	-
<i>Supervised Methods without Additional Features</i>								
<i>CopyRNN (Meng et al., 2017)</i>	0.174	0.331	0.413	0.217*	0.237*	0.210*	0.327	0.278
<i>DivGraphPointer (Sun et al., 2019)</i>	-	-	-	-	-	-	0.368	0.292
<i>Div-DGCN (Zhang et al., 2020)</i>	-	-	-	-	-	-	0.349	0.313
<i>SKE-Large-CLS (Mu et al., 2020)</i>	-	-	-	-	-	-	0.392	0.330
<i>ChunkKPE (Sun et al., 2020)</i>	0.283	0.486	0.581	0.355	0.373	0.324	0.408	0.337
<i>RankKPE (Sun et al., 2020)</i>	0.290	0.509	0.604	0.361	0.390	0.337	0.417	0.343
<i>JointKPE (Sun et al., 2020)</i>	0.291	0.511	0.605	0.364	0.391	0.338	0.419	0.344
<b><i>KIEMP</i></b>	<b>0.298</b>	<b>0.517</b>	<b>0.615</b>	<b>0.369</b>	<b>0.392</b>	<b>0.340</b>	<b>0.421</b>	<b>0.345</b>

Table 4: Performances of keyphrase extraction model on the *OpenKP* development set and the *KP20k* testing set. The best results of our model are highlighted in bold, and the best results of baselines are underlined. \* indicates these numbers are not included in the original paper and are estimated with Precision and Recall. The results of the baselines are reported in their corresponding papers.

target keyphrases and extracted keyphrases when determining the match of keyphrases and match of the identical word.

#### 4.4 Implementation Details

Implementation details of our proposed models are as follows. The maximum document length is 512 due to BERT limitations (Devlin et al., 2019), and documents are zero-padded or truncated to this length. The training used 6 GeForce RTX 2080 Ti GPUs and took about 31 hours and 77 hours for *OpenKP* and *KP20k* datasets respectively. Table 3 lists the parameters of our model. Furthermore, the model was implemented in Pytorch (Paszke et al., 2019) using the huggingface re-implementation of RoBERTa (Wolf et al., 2019).

## 5 Results and Analysis

This section investigates the performance of the proposed *KIEMP* on six widely-used benchmark datasets (*OpenKP*, *KP20k*, *Inspec*, *Krapivin*, *Nus*, and *Semeval*) from three facets. The first one demonstrates its superiority by comparing it with ten baselines in terms of several metrics. The second one is to verify the sensitivity of the concept dimension. The last one is to explicitly show the

keyphrase extraction results of *KIEMP* via two examples (two testing documents).

### 5.1 Overall Performance

The overall performance of different algorithms on two benchmarks (*OpenKP* and *KP20k*) is summarized in Table 4. We can see that the supervised methods outperform all the unsupervised algorithms (*TFIDF* and *TextRank*). This is not surprising since the supervised methods are trained end-to-end with supervised data. In all the supervised baselines, the methods using additional features are better than those without additional features. The reason is that the models with additional features are equal to encode keyphrases from multiple features perspectives. Therefore, it is helpful for the model to measure the importance of each keyphrase, thus improving the performance of the result of keyphrase extraction. Intuitively, this is the same as our proposed method. *KIEMP* considers the importance of keyphrases from multiple perspectives and fairly measures the importance of each keyphrase. But the difference is that we do not need additional features to assist. And in many practical applications of keyphrase extraction, there is no additional feature (i.e., visual features) infor-




Figure 2: Results of keyphrase extraction model on four testing sets (*Semeval*, *Inspec*, *Krapivin*, and *Nus*). The results of *JointKPE* are re-evaluated via the code which is provided by its corresponding paper.

mation to use in most cases. Compared with recent baselines (*ChunkKPE*, *RankKPE*, and *JointKPE*), *KIEMP* performs stably better on all metrics on both two datasets. These results demonstrate the benefits of estimating the importance of keyphrases from multiple perspectives simultaneously and the effectiveness of our multi-task learning strategy.

Furthermore, to verify the robustness of *KIEMP*, we also test the *KIEMP* trained with *KP20k* dataset on four widely-adopted keyphrase extraction data sets. It can be seen from Figure 2 that *KIEMP* is superior to the best baseline (*JointKPE*). We consider that this phenomenon comes from two benefits. One is that the high-level concepts captured by a deep latent variable model may contain topic and structure features. These features are essential information to evaluate the importance of phrases. Another one is that concept consistency matching uses a deep latent variable model to capture concepts. Here, the latent variable is characterized by a probability distribution over possible values rather than a fixed value, which can enforce the uncertainty of the model and further lead to robust phrase representation learning.

<i>Concept Dimension</i> ( $\mathbb{R}^c$ )	<i>OpenKP</i>		
	$R@1$	$R@3$	$R@5$
64	<b>0.298</b>	<b>0.517</b>	<b>0.615</b>
256	0.297	0.513	0.610
512	0.296	0.509	0.609
768	0.293	0.508	0.606

Table 5: Effectiveness of different dimensions of latent concept representation. The best results are highlighted in bold.

## 5.2 Sensitivity of the Concept Dimension

Here, we verify the effectiveness of using different concept dimensions. From Table 5, we can find that the increase of the dimension of latent concept representation has little effect on the result of keyphrase extraction. In contrast, the smaller the dimension, the better the result. Furthermore, in Table 4, the improvement of our proposed *KIEMP* model on the  $F_1@1$  evaluation metric is higher than the  $F_1@3$  and  $F_1@5$  evaluation metrics on the *KP20k* dataset. We consider the main reason is that our concept representation may capture the high-level conceptual information of phrases or documents, such as topics and structure information. Therefore, *KIEMP* with concept consistency matching module focuses more on extracting keyphrases closest to the main topic of the given document.

## 5.3 Case Study

To further illustrate the effectiveness of the proposed model, we present a case study on the results of the keyphrases extracted by different algorithms. Table 6 presents the results of *KIEMP* without concept consistency matching and *KIEMP*. From the first example, we can see that our *KIEMP* model is more inclined to extract keyphrases closer to the central semantics of the input document, which benefits from our concept consistency matching model. From the second example, we can see that the keyphrases extracted by *KIEMP* without concept consistency matching contain some redundant or meaningless phrases. The main reason may be that the *KIEMP* without concept consistency matching does not measure the importance of phrases from multiple perspectives, which leads to biased extraction. On the contrary, the keyphrases extracted by *KIEMP* are all around the main concepts of the example document, i.e., “leadership”. It further demonstrates the effectiveness of our proposed model.

## 6 Conclusions and Future Work

A new keyphrase importance estimation from the multiple perspectives approach is proposed to estimate the importance of keyphrase. Benefiting from the designed syntactic accuracy chunking, information saliency ranking, and concept consistency matching modules, *KIEMP* can fairly extract keyphrases. A series of experiments have demonstrated that *KIEMP* outperformed the existing state-of-the-art keyphrase extraction methods. In the



<p>(A) <i>Part of the Input Document:</i></p> <p>The Great Plateau is a large region of land that is secluded from other parts of Hyrule, as its steep slopes prevent anyone from traveling to and from it without special equipment, such as the Paraglider. The only active inhabitant is the Old Man, a mysterious ... (URL: <a href="https://zelda.gamepedia.com/Great_Plateau">https://zelda.gamepedia.com/Great_Plateau</a>)</p> <p><i>Target Keyphrase:</i>  (1) <b>great plateau</b> ; (2) <b>breath of the wild</b> ; (3) <b>hyrule</b></p>
<p><i>KIEMP without concept consistency matching:</i>  (1) <b>great plateau</b> ; (2) <b>hyrule</b> ; (3) <b>breath of the wild</b> ; (4) paraglider ; (5) zelda</p> <p><i>KIEMP:</i>  (1) <b>great plateau</b> ; (2) <b>breath of the wild</b> ; (3) <b>hyrule</b> ; (4) paraglider ; (5) starting region</p>
<p>(B) <i>Part of the Input Document:</i></p> <p>Transformational leaders also depend on visionary leadership to win over followers, but they have an added focus on employee development. For example, a transformational leader might explain how her plan for the future serves her employees' interests and how she will support them through the changes ... (URL: <a href="https://yourbusiness.azcentral.com/managers-different-leadership-styles-motivate-teams-8481.html">https://yourbusiness.azcentral.com/managers-different-leadership-styles-motivate-teams-8481.html</a>)</p> <p><i>Target Keyphrase:</i>  (1) <b>managers</b> ; (2) <b>leadership</b> ; (3) <b>teams</b></p>
<p><i>KIEMP without concept consistency matching:</i>  (1) motivating ; (2) motivate ; (3) charismatic leadership ; (4) transformational leadership ; (5) employee development</p> <p><i>KIEMP:</i>  (1) leadership styles ; (2) <b>managers</b> ; (3) charismatic leadership ; (4) transformational leadership ; (5) <b>leadership</b></p>

Table 6: Example of keyphrase extraction results (selected from the *OpenKP* dataset). Phrases in red and bold are target keyphrases predicted by the different models (*KIEMP* without concept consistency matching and *KIEMP*).

future, it will be interesting to introduce an adaptive approach in *KIEMP* to filter the meaningless phrases.

## 7 Acknowledgments

This work was supported in part by the National Key Research and Development Program of China under Grant 2020AAA0106800; the National Science Foundation of China under Grant 61822601 and 61773050; the Beijing Natural Science Foundation under Grant Z180006; The Fundamental Research Funds for the Central Universities (2019JBZ110).

## References

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). In *NAACL-HLT*, pages 4171–4186. Association for Computational Linguistics.
- Maria P. Grineva, Maxim N. Grinev, and Dmitry Lizorkin. 2009. [Extracting key terms from noisy and multitheme documents](#). In *WWW*, pages 661–670. ACM.
- Kazi Saidul Hasan and Vincent Ng. 2014. [Automatic keyphrase extraction: A survey of the state of the art](#). In *ACL (1)*, pages 1262–1273. The Association for Computer Linguistics.
- Anette Hulth. 2003. [Improved automatic keyword extraction given more linguistic knowledge](#). In *EMNLP*.
- Anette Hulth. 2004. [Enhancing linguistically oriented automatic keyword extraction](#). In *HLT-NAACL (Short Papers)*. The Association for Computational Linguistics.
- Xin Jiang, Yunhua Hu, and Hang Li. 2009. [A ranking approach to keyphrase extraction](#). In *SIGIR*, pages 756–757. ACM.
- Karen Spärck Jones. 2004. [A statistical interpretation of term specificity and its application in retrieval](#). *J. Documentation*, 60(5):493–502.
- Su Nam Kim, Olena Medelyan, Min-Yen Kan, and Timothy Baldwin. 2010. [Semeval-2010 task 5 : Automatic keyphrase extraction from scientific articles](#). In *SemEval@ACL*, pages 21–26. The Association for Computer Linguistics.
- M. Krapivin and M. Marchese. 2009. Large dataset for keyphrase extraction.
- Feifan Liu, Deana Pennell, Fei Liu, and Yang Liu. 2009a. [Unsupervised approaches for automatic keyword extraction using meeting transcripts](#). In *HLT-NAACL*, pages 620–628. The Association for Computational Linguistics.

- Yang Liu and Mirella Lapata. 2019. [Text summarization with pretrained encoders](#). In *EMNLP/IJCNLP (1)*, pages 3728–3738. 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](#). *CoRR*, abs/1907.11692.
- Zhiyuan Liu, Wenyi Huang, Yabin Zheng, and Maosong Sun. 2010. [Automatic keyphrase extraction via topic decomposition](#). In *EMNLP*, pages 366–376. ACL.
- Zhiyuan Liu, Peng Li, Yabin Zheng, and Maosong Sun. 2009b. [Clustering to find exemplar terms for keyphrase extraction](#). In *EMNLP*, pages 257–266. ACL.
- O. Medelyan, E. Frank, and I. H. Witten. 2009. [Human-competitive tagging using automatic keyphrase extraction](#). In *Internat. Conference of Empirical Methods in Natural Language Processing, EMNLP-2009*.
- Rui Meng, Sanqiang Zhao, Shuguang Han, Daqing He, Peter Brusilovsky, and Yu Chi. 2017. [Deep keyphrase generation](#). In *ACL*, pages 582–592. Association for Computational Linguistics.
- Rada Mihalcea and Paul Tarau. 2004. [Textrank: Bringing order into text](#). In *EMNLP*, pages 404–411. ACL.
- Funan Mu, Zhenting Yu, Lifeng Wang, Yequan Wang, Qingyu Yin, Yibo Sun, Liqun Liu, Teng Ma, Jing Tang, and Xing Zhou. 2020. [Keyphrase extraction with span-based feature representations](#). *CoRR*, abs/2002.05407.
- Chau Q. Nguyen and Tuoi T. Phan. 2009. [An ontology-based approach for key phrase extraction](#). In *ACL/IJCNLP (Short Papers)*, pages 181–184. The Association for Computer Linguistics.
- Thuy Dung Nguyen and Min-Yen Kan. 2007. [Keyphrase extraction in scientific publications](#). In *ICADL*, volume 4822 of *Lecture Notes in Computer Science*, pages 317–326. Springer.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *NeurIPS*, pages 8024–8035.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *NAACL-HLT*, pages 2227–2237. Association for Computational Linguistics.
- M.F. Porter. 2006. An algorithm for suffix stripping. *Program: Electronic Library and Information Systems*, 40(3):211–218.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. 2014. [Stochastic backpropagation and approximate inference in deep generative models](#). Cite arxiv:1401.4082Comment: Appears In Proceedings of the 31st International Conference on Machine Learning (ICML), JMLR: W&CP volume 32, 2014.
- Shalena Srna, Rom Y. Schrift, and Gal Zauberan. 2018. [The illusion of multitasking and its positive effect on performance](#). *Psychological Science*, 29(12):1942–1955.
- Si Sun, Chenyan Xiong, Zhenghao Liu, Zhiyuan Liu, and Jie Bao. 2020. [Joint keyphrase chunking and salience ranking with bert](#). *CoRR*, abs/2004.13639.
- Zhiqing Sun, Jian Tang, Pan Du, Zhi-Hong Deng, and Jian-Yun Nie. 2019. [Divgraphpointer: A graph pointer network for extracting diverse keyphrases](#). In *SIGIR*, pages 755–764. ACM.
- Takashi Tomokiyo and Matthew Hurst. 2003. [A language model approach to keyphrase extraction](#). pages 33–40. Association for Computational Linguistics.
- Peter D. Turney. 2000. [Learning algorithms for keyphrase extraction](#). *Inf. Retr.*, 2(4):303–336.
- Peter D. Turney. 2002. [Learning to extract keyphrases from text](#). *CoRR*, cs.LG/0212013.
- Xiaojun Wan and Jianguo Xiao. 2008. [Collabrank: Towards a collaborative approach to single-document keyphrase extraction](#). In *COLING*, pages 969–976.
- Yansen Wang, Zhen Fan, and Carolyn Penstein Rosé. 2020. [Incorporating multimodal information in open-domain web keyphrase extraction](#). In *EMNLP (1)*, pages 1790–1800. Association for Computational Linguistics.
- Yue Wang, Jing Li, Hou Pong Chan, Irwin King, Michael R. Lyu, and Shuming Shi. 2019. [Topic-aware neural keyphrase generation for social media language](#). In *ACL (1)*, pages 2516–2526. Association for Computational Linguistics.
- Ian H. Witten, Gordon W. Paynter, Eibe Frank, Carl Gutwin, and Craig G. Nevill-Manning. 1999. [Kea: Practical automatic keyphrase extraction](#). In *ACM DL*, pages 254–255. ACM.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *CoRR*, abs/1910.03771.
- Lee Xiong, Chuan Hu, Chenyan Xiong, Daniel Campos, and Arnold Overwijk. 2019. [Open domain web keyphrase extraction beyond language modeling](#). In *EMNLP/IJCNLP (1)*, pages 5174–5183. Association for Computational Linguistics.

Hongyuan Zha. 2002. [Generic summarization and keyphrase extraction using mutual reinforcement principle and sentence clustering](#). In *SIGIR*, pages 113–120. ACM.

Haoyu Zhang, Dingkun Long, Guangwei Xu, Pengjun Xie, Fei Huang, and Ji Wang. 2020. [Keyphrase extraction with dynamic graph convolutional networks and diversified inference](#). *CoRR*, abs/2010.12828.