

# Enhancing Phrase Representation by Information Bottleneck Guided Text Diffusion Process for Keyphrase Extraction

Yuanzhen Luo<sup>1†</sup>, Qingyu Zhou<sup>2\*</sup>, Feng Zhou<sup>2</sup>

<sup>1</sup>China University of Petroleum, Beijing

<sup>2</sup>OPPO Research Institute

strugglingluo@gmail.com, qyzhgm@gmail.com, zhoufeng1@oppo.com

## Abstract

Keyphrase extraction (KPE) is an important task in Natural Language Processing for many scenarios, which aims to extract keyphrases that are present in a given document. Many existing supervised methods treat KPE as sequential labeling, span-level classification, or generative tasks. However, these methods lack the ability to utilize the reference keyphrase information during extraction process, which may result in inferior results. In this study, we propose Diff-KPE, which leverages the supervised Variational Information Bottleneck (VIB) to guide the text diffusion process for generating enhanced keyphrase representations. Diff-KPE first generates the desired keyphrase embeddings conditioned on the entire document and then injects the generated keyphrase embeddings into each phrase representation. A ranking network and VIB are then optimized together with rank loss and classification loss, respectively. This design of Diff-KPE allows us to rank each candidate phrase by utilizing both the information of keyphrases and the document. Experiments show that Diff-KPE outperforms most of existing KPE methods on a large open domain keyphrase extraction benchmark, OpenKP, and a scientific domain dataset, KP20K.

**Keywords:** Keyphrase Extraction, Diffusion, Information Bottleneck

## 1. Introduction

Keyphrase extraction (KPE) aims to extract several *present* keyphrases from a document that can highly summarize the given document, which is helpful for many applications such as text summarization and information retrieval.

Many neural keyphrase extraction models formulate KPE as a token-level sequence labeling problem by predicting a single label for each token (Sahrawat et al., 2020; Alzaidy et al., 2019; Luan et al., 2017). To use the phrase-level semantic information, some methods (Zhang et al., 2016; Xiong et al., 2019; Mu et al., 2020; Wang et al., 2020; Sun et al., 2021) modeling KPE as a phrase classification task by assigning labels to each text span.

Different from the above methods, recent KPE models directly learn to rank each phrase (Mu et al., 2020; Song et al., 2021; Sun et al., 2021; Song et al., 2022). These methods mainly include two processes: candidate phrase representation construction and keyphrase importance ranking. In particular, candidate phrase representations are extracted from the token embeddings produced by pre-trained language models such as BERT (Devlin et al., 2018), and keyphrase importance ranking usually predict a score for each candidate phrase representation and then use margin loss to sort the score of positive candidates ahead of negative ones. Since the candidate phrase representation

is important for the model to score them, there are several ways to extract phrase representation: (Mu et al., 2020) and (Sun et al., 2021; Song et al., 2021) develop Bi-LSTM and CNN to further capture the local-aware features of phrase, respectively. HyperMatch (Song et al., 2022) extract phrase representation in hyperbolic space.

Although these methods have achieved great success in many KPE benchmarks, we point out that their candidate phrase representation still lacks the utilization of reference keyphrases information. This is inspired by the intuition that how human extracts candidate phrase: They will first review the whole document and summarize a few *vague* keyphrases in mind, and then take the candidate phrase and vague keyphrases both into consideration to make the extraction decision. To achieve this process in the neural model, however, the challenge is how to generate the vague reference keyphrase information during inference time.

To address the above issue, we propose Diff-KPE, a novel diffusion-based KPE model. We first use the diffusion model to generate a list of vague keyphrase information by recovering reference keyphrase embeddings conditioned on the whole document, then we enhance the phrase representation by injecting the vague keyphrase embeddings into each of them. To rank candidate phrases, we apply a ranking network to rank each enhanced phrase representation. By doing this, we can extract desired top  $k$  keyphrases from the ranked list of phrases. In addition, we introduce a supervised Variational Information Bottleneck (VIB) to optimize a classification loss for each

<sup>†</sup>Work done during internship at OPPO Research Institute.

\*Corresponding author.

phrase. Supervised VIB aims to preserve the information about the target classes in the latent space while filtering out irrelevant information from the input phrase representation (Tishby et al., 2000), which helps the learning process for the vague keyphrase embedding. Multitask learning of supervised VIB can guide the model to generate informative phrase representations, thereby improving the performance of the ranking network. Overall, Diff-KPE incorporates these modules by simultaneously training these components.

Empowered by the architecture design of Diff-KPE, it exhibits the following three advantages. First, the diffusion model enables the injection of vague keyphrase information into each phrase representation, even during inference, thereby giving the model the ability to utilize keyphrase information. Second, the ranking network ranks each phrase, allowing us to flexibly extract the top  $k$  candidate keyphrases. Finally, the introduced supervised VIB guides the model to generate informative phrase representations, resulting in an improvement in ranking performance. We demonstrate the importance of each component in our experiments. In summary, the main contributions of this paper are as follows:

- We propose Diff-KPE, a diffusion-based KPE model. To the best of our knowledge, this is the first attempt to use the diffusion model for the KPE task.
- By incorporating the diffusion model, ranking network, and VIB into one system, we empower Diff-KPE to utilize the information of keyphrases and the document to extract candidate keyphrases.
- Experimental results show that Diff-KPE outperforms most of existing KPE approaches on two large keyphrase extraction benchmarks, OpenKP, and KP20K. Additionally, Diff-KPE demonstrates a more robust performance on the other five small scientific datasets.

## 2. Related Work

### 2.1. Keyphrase Extraction

Automatic KeyPhrase Extraction (KPE) aims to extract a set of important and topical phrases from a given document, which can then be used in different tasks such as summarization (Li et al., 2020), problem solving (Huang et al., 2017, 2018b,a), generation tasks (Zhou and Huang, 2019; Li et al., 2021), and so on.

Existing KPE technologies can be categorized as unsupervised and supervised methods. Unsupervised methods are mainly based on statistical information (El-Beltagy and Rafea, 2009; Florescu

and Caragea, 2017a; Campos et al., 2018), embedding features (Mahata et al., 2018; Sun et al., 2020; Liang et al., 2021; Zhang et al., 2021; Ding and Luo, 2021), and graph-based ranking algorithms (Mihalcea and Tarau, 2004; Florescu and Caragea, 2017b; Boudin, 2018). Supervised methods commonly formulate KPE as sequence tagging approaches (Sahrawat et al., 2020; Alzaidy et al., 2019; Kulkarni et al., 2022), span-level classification (Zhang et al., 2016; Xiong et al., 2019; Mu et al., 2020; Sun et al., 2021) or ranking (Mu et al., 2020; Song et al., 2021; Sun et al., 2021; Song et al., 2022), or generative tasks (Meng et al., 2017; Chen et al., 2018; Yuan et al., 2018; Kulkarni et al., 2022). Although supervised KPE methods require a lot of annotated data, their performance is significantly superior to unsupervised methods in many KPE benchmarks (Sun et al., 2021; Meng et al., 2017).

Recently, some works have focused on constructing a zero-shot keyphrase extractor by prompting pre-trained large language models (LLMs). For example, (Song et al., 2023) verified the performance of ChatGPT (OpenAI, 2022) and ChatGLM-6b (Zeng et al., 2022) for the KPE task and found that they still have a lot of room for improvement in the KPE task compared to existing SOTA supervised models. Similar results can also be observed in (Martínez-Cruz et al., 2023).

### 2.2. Diffusion Models for Text

Diffusion models have been applied in many continuous domain generations like image, video, and audio (Kong et al., 2020; Rombach et al., 2022; Ho et al., 2022; Yang et al., 2022). Recently, there are some works focused on applying the diffusion model to discrete text data. They usually generate continuous representations for the desired texts/words. For example, Diffusion-LM (Li et al., 2022) first attempts to develop a continuous diffusion model to generate text by embedding rounding step. Following the work of Diffusion-LM, DiffuSeq (Gong et al., 2022) and SeqDiffuSeq (Yuan et al., 2022) designed a diffusion-based sequence-to-sequence model for the text generation task. To adapt the diffusion model to longer sequence generation, (Zhang et al., 2023) proposed a sentence-level diffusion generation model for summary tasks, which directly generates sentence-level embeddings and matches from embeddings back to the original text.

Contrary to previous works, we apply the diffusion model to KPE, a phrase-level extraction task. In order to take the keyphrase information into consideration during extraction, we directly inject the keyphrase information generated by the diffusion model into each phrase representation.

### 2.3. Variational Information Bottleneck in NLP

Variational Information Bottleneck (VIB) is one of a group of Information Bottleneck (IB) methods. It aims to find compact representations of data that preserve the most relevant information while filtering out irrelevant or redundant information (Tishby et al., 2000).

There are lots of studies that apply VIB to many NLP tasks. For example, (Li and Eisner, 2019) used VIB for parsing, and (West et al., 2019) used it for text summarization. Recently, VIB was also used in Named Entity Recognition (NER) (Wang et al., 2022; Nguyen et al., 2023), text classification (Zhang et al., 2022), machine translation (Ormazabal et al., 2022) and so on.

## 3. Methodology

In this section, we present the detailed design of our Keyphrase Extraction (KPE) model, named Diff-KPE. An overview of Diff-KPE is depicted in Figure 1. Given document  $D = \{w_1, w_2, \dots, w_n\}$ , we start by enumerating all possible phrase representations and reference keyphrase embeddings. The Diffusion module is then employed to reconstruct the keyphrase embeddings and inject them into each phrase representation. Additionally, we incorporate a supervised Variational Information Bottleneck (VIB) for phrase classification and a ranking network for ranking purposes. These components work together to enhance the performance of our KPE model.

### 3.1. Phrase Representation

To enumerate and encode all the possible phrase representations, we first use pre-trained language model BERT (Devlin et al., 2018) to encode document  $D = \{w_1, w_2, \dots, w_n\}$ , producing contextual word embeddings  $\mathbf{E} = \{e_1, e_2, \dots, e_n\}$ . The word embeddings are then integrated into phrase representations using a set of Convolutional Neural Networks (CNNs):

$$\mathbf{s}_i^k = \text{CNN}^k(e_i, e_{i+1}, \dots, e_{i+n-1}) \quad (1)$$

where  $1 \leq k \leq N$  and  $N$  represents the pre-defined maximum length of phrase. The  $i$ th  $k$ -gram phrase representation  $\mathbf{s}_i^k$  is calculated by its corresponding  $\text{CNN}^k$ .

### 3.2. Keyphrase Embeddings Generation

In order to inject reference keyphrases information into each phrase representation, we use a continuous diffusion module to generate desired keyphrase embeddings.

#### 3.2.1. Input Encoding

To allow the diffusion module to generate desired keyphrase embeddings conditioned on the whole document, we first use another BERT model to obtain initial document and keyphrases embeddings. Refer to  $m$  keyphrases and document embedding as  $\mathbf{E}^{kp} = \{e_i^{kp}\}_i^m$  and  $e^D$ , the input encoding of the diffusion module is formatted as:

$$\begin{aligned} \mathbf{H}^{\text{in}} &= \mathbf{h}^D \parallel \mathbf{H}^{kp} \\ &= \text{TransformerEncoder}(e^D \parallel \mathbf{E}^{kp}) \end{aligned} \quad (2)$$

where  $\mathbf{H}^{kp} = \{\mathbf{h}_i^{kp}\}_i^m$  and  $\mathbf{h}^D$  are the latent embedding of the document and  $m$  keyphrases, **TransformerEncoder** is a stacked Transformer encoder which embeds the input vector into latent space, and  $e^D$  is the document embedding, i.e., the [CLS] token embedding in BERT model, and  $\parallel$  indicates concatenation operation. Such input encoding enables our continuous diffusion module to generate desired keyphrase embeddings conditional to the current document embeddings  $e^D$ .

#### 3.2.2. Diffusion Generation Process

Once the input encoding  $\mathbf{H}^{\text{in}}$  is obtained, the diffusion model aims to perturb  $\mathbf{H}^{\text{in}}$  gradually and then recover the original  $\mathbf{H}^{\text{in}}$  by learning a reverse process. To achieve this, a one-step Markov transition  $q(\mathbf{x}_0 | \mathbf{H}^{\text{in}})$  is performed to obtain the initial state  $\mathbf{x}_0$ :

$$\begin{aligned} \mathbf{x}_0 &= \mathbf{x}_0^D \parallel \mathbf{x}_0^{kp} \\ &\sim \mathcal{N}(\mathbf{H}^{\text{in}}, \beta_0 \mathbf{I}) \end{aligned} \quad (3)$$

where  $\beta_t \in (0, 1)$  adjusts the scale of the variance,  $\mathbf{x}_0^D \sim \mathcal{N}(\mathbf{h}^D, \beta_0 \mathbf{I})$  and  $\mathbf{x}_0^{kp} \sim \mathcal{N}(\mathbf{H}^{kp}, \beta_0 \mathbf{I})$  are the latent document embedding and keyphrase embeddings, respectively. We then start the forward process by gradually adding Gaussian noise to the latent keyphrase embeddings  $\mathbf{x}_t^{kp}$ . Following the previous work (Zhang et al., 2023), we keep the latent document embedding  $\mathbf{x}_0^D$  unchanged, so that the diffusion module can generate keyphrase embeddings condition to the source document. Formally, at step  $t$  of the forward process  $q(\mathbf{x}_t | \mathbf{x}_{t-1})$ , the noised latent embedding is  $\mathbf{x}_t$ :

$$\mathbf{x}_t = \mathbf{x}_0^D \parallel \mathcal{N}(\mathbf{x}_t^{kp}; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}^{kp}, \beta_t \mathbf{I}) \quad (4)$$

where  $t \in \{1, 2, \dots, T\}$  for a total of  $T$  diffusion steps. For more details about the diffusion generation process, please refer to (Sohl-Dickstein et al., 2015).

After adding the noise gradually at a specific time step  $t$  (usually randomly choose between  $[1, T]$ ), the backward process is performed to recover the keyphrase embeddings  $\mathbf{x}_t^{kp}$  by removing the noised. We use another stacked Transformer

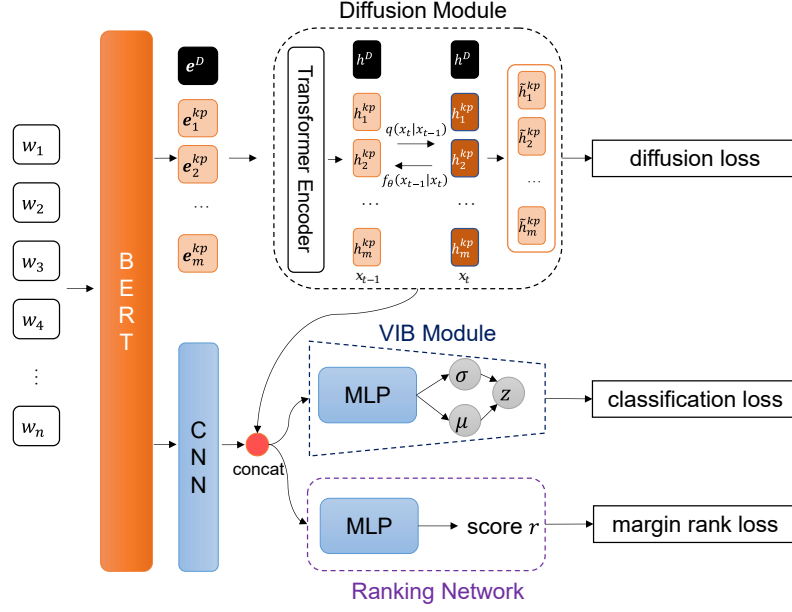


Figure 1: Diff-KPE is jointly trained with a continuous diffusion module, a variational information bottleneck, and a rank network. The black dashed box is the diffusion module, the blue dashed box is the VIB module and the purple dashed box is the rank network.

encoder model  $f_\theta$  to conduct this backward process to recover the original input encoding  $\mathbf{H}^{kp}$ :

$$\tilde{\mathbf{H}}^{kp} = f_\theta(\mathbf{x}_t^{kp}, t) \quad (5)$$

where  $f_\theta(\mathbf{x}_t^{kp}, t)$  is the stacked Transformer network to reconstruct  $\mathbf{H}^{kp}$  at time step  $t$ .

Since the main objective of the diffusion generation module is to reconstruct the original input encoding, the objective loss of continuous diffusion module can be defined by:

$$\mathcal{L}_{diff} = \sum_{t=1}^T \|\mathbf{H}^{kp} - f_\theta(\mathbf{x}_t^{kp}, t)\|^2 + \mathcal{R}(\mathbf{x}_0) \quad (6)$$

where  $\mathcal{R}(\mathbf{x}_0)$  is a regularization term for  $\mathbf{x}_0$ .

### 3.3. Keyphrase Ranking

After the diffusion generation process, the generated keyphrase embeddings  $\tilde{\mathbf{H}}^{kp}$  are concatenated into each phrase representation  $\tilde{s}_i^k$ . This aims to inject the information from keyphrases into each phrase, resulting in performance improvement of keyphrase ranking. Specifically, formulate the final phrase representation as:

$$\tilde{s}_i^k = s_i^k \parallel \text{flat}(\tilde{\mathbf{H}}^{kp}) \quad (7)$$

where  $\text{flat}(\mathbf{x})$  means that  $\mathbf{x}$  is flattened to a vector. Equation 7 means that the final phrase representation not only contains the original phrase representation but also all the reconstructed keyphrase information.

For training the model to rank each phrase, we introduce a contrastive rank loss. Following the previous work (Sun et al., 2021), we first take a feedforward layer to project the input representation  $\tilde{s}_i^k$  to a scalar score:

$$r(\tilde{s}_i^k) = \text{FeedForward}(\tilde{s}_i^k) \quad (8)$$

Then the margin rank loss is introduced to learn to rank keyphrase  $\tilde{s}_{s+}$  ahead of non-keyphrase  $\tilde{s}_{s-}$  for the given document  $D$ :

$$\mathcal{L}_{rank} = \sum_{\tilde{s}_{s+}, \tilde{s}_{s-} \in D} \max(0, 1 - r(\tilde{s}_{s+}) + r(\tilde{s}_{s-})) \quad (9)$$

### 3.4. Keyphrase Classification

Combining the keyphrase classification task during training can enhance the phraseness measurement of the phrase (Sun et al., 2021; Song et al., 2021). Similar to previous work (Xiong et al., 2019; Sun et al., 2021; Song et al., 2021), we introduce a classification loss for each final phrase representation for multi-task learning. We found that the use of supervised VIB substantially improves the ranking performance (See Ablation Study). Supervised VIB aims to preserve the information about the target classes in the latent space while filtering out irrelevant information from the input (Voloshynovskiy et al., 2019). Given the final phrase representation  $\tilde{s}_i^k$ , the supervised VIB first compresses the input to a latent variable  $z \sim q_{\phi_1}(z|\tilde{s}_i^k)$ . We apply two linear layers to construct the parameters  $q$  using

the following equations:

$$\begin{aligned}\mu &= \mathbf{W}_\mu \tilde{\mathbf{s}}_i^k + \mathbf{b}_\mu \\ \sigma^2 &= \mathbf{W}_\sigma \tilde{\mathbf{s}}_i^k + \mathbf{b}_\sigma\end{aligned}\quad (10)$$

where  $\mu$  and  $\sigma$  are the parameters of a multivariate Gaussian, representing the latent feature space of the phrase;  $\mathbf{W}$  and  $\mathbf{b}$  are weights and biases of the linear layer, respectively. The posterior distribution  $z \sim q_{\phi_1}(z|\tilde{\mathbf{s}}_i^k)$  is approximated via reparameterisation trick (Kingma and Welling, 2013):

$$z = \mu + \sigma \epsilon, \text{ where } \epsilon \sim \mathcal{N}(0, 1) \quad (11)$$

Since the main objective of VIB is to preserve target class information while filtering out irrelevant information from the input, the objective loss function for the supervised VIB is based on classification loss and compression loss. Denoted by  $y$  as the true label of the input phrase, the objective loss of the supervised VIB is defined as:

$$\begin{aligned}\mathcal{L}_{vib}(\phi) &= \mathbb{E}_z[-\log p_{\phi_2}(y|z)] \\ &+ \alpha \mathbb{E}_{\tilde{\mathbf{s}}_i^k}[D_{KL}(q_{\phi_1}(z|\tilde{\mathbf{s}}_i^k), pr(z))]\end{aligned}\quad (12)$$

where  $pr(z)$  is an estimate of the prior probability  $q_{\phi_1}(z)$ ,  $\alpha$  range in  $[0, 1]$ ,  $\phi$  is the neural network parameters, and  $D_{KL}$  is the Kullback-Leibler divergence. We use a multi-layer perceptron with one linear layer and softmax function to calculate  $p_{\phi_2}(y|z)$ . Note that Equation 12 can be approximated by the Monte Carlo sampling method with sample size  $M$ .

### 3.5. Optimization and Inference

We jointly optimize the diffusion module, ranking network, and supervised VIB end-to-end. Specifically, the overall training objective loss can be represented as:

$$\mathcal{L} = \mathcal{L}_{dif} + \mathcal{L}_{vib} + \mathcal{L}_{rank} \quad (13)$$

For inference, the Transformer encoder first obtains the initial document embeddings  $\mathbf{h}^D$ , and then the one-step Markov  $q(\mathbf{x}_0^D|\mathbf{h}^D)$  is performed. To construct the noise keyphrase embedding  $\mathbf{x}_T^{kp}$ , we random sample  $m$  Gaussian noise embeddings such that  $\mathbf{x}_T^{kp} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . Then the reverse process is applied to remove the Gaussian noise of  $\mathbf{x}_T = \mathbf{x}_0^D|\mathbf{x}_T^{kp}$  iteratively and get the output keyphrase embeddings  $\tilde{\mathbf{H}}^{kp} = [\tilde{\mathbf{h}}_1^{kp}, \tilde{\mathbf{h}}_2^{kp}, \dots, \tilde{\mathbf{h}}_m^{kp}]$ . After that, each original phrase representation  $\mathbf{s}_i^k$  is concatenated to the flattened keyphrase embeddings  $\tilde{\mathbf{H}}^{kp}$  and input to the ranking network to obtain the final score for each phrase.

## 4. Experiments

### 4.1. Datasets

In this paper, we use seven KPE benchmark datasets in our experiments.

- **OpenKP** (Xiong et al., 2019) consists of around 150K web documents from the Bing search engine. We follow its official split of training (134K), development (6.6K), and testing (6.6K) sets. Each document in OpenKP was labeled with 1-3 keyphrases by expert annotators.
- **KP20K** (Meng et al., 2017) consists of a large amount of high-quality scientific metadata in the computer science domain from various online digital libraries (Meng et al., 2017). We follow the original partition of training (528K), development (20K), and testing (20K) set.
- **SemEval-2010** (Kim et al., 2013) contains 244 scientific documents. The official split of 100 testing documents is used for testing in our experiments.
- **SemEval-2017** (Augenstein et al., 2017) contains 400 scientific documents. The official split of 100 testing documents is used for testing in our experiments.
- **Nus** (Nguyen and Kan, 2007) contains 211 scholarly documents. We treat all 211 documents as testing data.
- **Inspec** (Hulth, 2003) contains 2000 paper abstracts. We use the original 500 testing papers and their corresponding controlled (extractive) keyphrases for testing.
- **Krapivin** (Krapivin et al., 2009) contains 2305 papers from scientific papers in ACM. We treat all 2305 papers as testing data.

Dataset	Avg. Doc Len.	Avg. KP Len.	Avg. # KP	% up to 5-gram
OpenKP	1212.3	2.0	2.2	99.2%
KP20k	169.3	1.9	3.5	99.8%
SemEval-2010	9664.2	2.0	9.5	99.8%
SemEval-2017	190.6	2.3	11.3	97.9%
Nus	8707.4	1.9	8.0	99.8%
Inspec	138.9	2.2	6.4	99.8%
Krapivin	9354.1	1.9	3.8	99.9%

Table 1: Statistics of benchmark datasets, including the average length of the document, the average length of the keyphrase, the average number of extractive keyphrases, and the percentage of keyphrases with a maximum length of 5.

Note that in order to verify the robustness of our model, we test the model trained with KP20K on

the testing data of SemEval-2010, SemEval-2017, Nus, Inspec, and Krapivin. For all datasets, only the *present* keyphrases are used for training and testing. The statistics of the training set of OpenKP and KP20k are shown in Table 1.

## 4.2. Baselines

To keep consistent with previous work (Meng et al., 2017; Xiong et al., 2019; Mu et al., 2020; Sun et al., 2021), we compare our model with two categories of KPE methods: Traditional KPE baselines and Neural KPE baselines.

Traditional KPE baselines consist of two popular unsupervised KPE methods, statistical feature-based method TF-IDF (Sparck Jones, 1972) and graph-based method TextRank (Mihalcea and Tarau, 2004), and two feature-based KPE systems PROD (Xiong et al., 2019) and Maui (Medelyan et al., 2009).

Neural KPE baselines consist of a sequence-to-sequence generation-based model named Copy-RNN (Meng et al., 2017). Previous state-of-the-art method on OpenKP and KP20K, KIEMP (Song et al., 2021) incorporating multiple perspectives estimation for phrase ranking. Another strong baseline, JointKPE (Sun et al., 2021), including its two variants ChunkKPE and RankKPE are reproduced according to their open-source code<sup>1</sup>. HyperMatch (Song et al., 2022), a new matching method for extracting keyphrase in the hyperbolic space. Two phrase-level classification-based models named SKE-Base-Cls (Mu et al., 2020) and BLING-KPE (Xiong et al., 2019). We also compare our model with BERT-based span extraction and sequence tagging methods, both of which come from the implementation of (Sun et al., 2021). Note that since both of KeyBart and KBIR (Kulkarni et al., 2022) are pre-trained with a well-defined pretraining strategy specifically on RoBERTa-large (Liu et al., 2019), we do not compare our Diff-KPE with them for a fair comparison.

In addition, we also add the results of two large language models (LLMs) with zero-shot settings: ChatGLM2-6b (Zeng et al., 2022) and ChatGPT<sup>2</sup> (OpenAI, 2022). To restrict the output format of ChatGPT, we design the following prompt template:

```
[Instruction]
```

```
Please extract 1 to 15 keyphrases
from the given document. Your ex-
tracted keyphrases should reason-
ably represent the topic of the docu-
ment and must appear in the original
text. You must give the keyphrases
by strictly following this format:
```

```
"[extracted keyphrases]", for ex-
ample: "[machine learning, neural
networks, NLP]"
```

```
[Document]
```

```
{document}
```

It should be noted that designing more complex prompts may improve the performance of LLMs, which is beyond the scope of this paper.

## 4.3. Evaluation Metrics

We use Recall (R), and F-measure (F1) of the top  $K$  predicted keyphrases for evaluating the performance of the KPE models. Following the prior research (Meng et al., 2017; Xiong et al., 2019), we utilize  $K = \{1, 3, 5\}$  on OpenKP and  $K = \{5, 10\}$  on others. When determining the exact match of keyphrases, we first lowercase the candidate keyphrases and reference keyphrases, and then we apply Porter Stemmer (Porter, 1980) to both of them.

## 4.4. Implementation details

We truncate or zero-pad each document due to the input length limitations (512 tokens). We use the base version of BERT to generate initial word embeddings. We also use the base version of Sentence-BERT (Reimers and Gurevych, 2019) to generate initial fixed phrase embeddings for the diffusion module. The maximum length of k-gram is set to  $N = 5$  for all datasets. The maximum diffusion time steps  $T$  is set to 100,  $\alpha = 2.8e - 6$ . The hidden size and number of layer in Transformer encoder in the diffusion module are set to 8 and 6 respectively. The latent dimension of the VIB model is set to 128. Sample size  $M = 5$ . We optimize Diff-KPE using AdamW with  $5e-5$  learning rate, 0.1 warm-up proportion, and 32 batch size. The training used 8 NVIDIA Tesla V100 GPUs and took about 20 hours on 5 epochs. During training Diff-KPE, we also set a simple early stop strategy such that the model would stop training if the validation performance (F1@3 for OpenKP, F1@5 for KP20K) does not improve after 5 times consecutive evaluations (We evaluate the model every 200 optimization steps), and we select the model with the best validation performance. We run our model with 5 different random seeds and report their average score.

## 5. Results and Analysis

In this section, we present the evaluation results of the proposed Diff-KPE on seven widely-used benchmark datasets (OpenKP, KP20k, SemEval-2010, SemEval-2017, Nus, Inspec, Krapivin).

<sup>1</sup><https://github.com/thunlp/BERT-KPE>

<sup>2</sup>GPT-3.5-turbo, version: gpt-3.5-turbo-0125

Model	OpenKP						KP20k		Average
	F1@1	R@1	F1@3	R@3	F1@5	R@5	F1@5	F1@10	
<b>Traditional KPE</b>									
TF-IDF <sup>†</sup> (Sun et al., 2021)	19.6	15	22.3	28.4	19.6	34.7	10.8	13.4	20.4
TextRank <sup>†</sup> (Sun et al., 2021)	5.4	4.1	7.6	9.8	7.9	14.2	18.0	15.0	10.2
Maui <sup>†</sup> (Mu et al., 2020)	-	-	-	-	-	-	27.3	24.0	-
PROD <sup>†</sup> (Xiong et al., 2019)	24.5	18.8	23.6	29.9	18.8	33.1	-	-	-
<b>Neural KPE</b>									
CopyRNN <sup>†</sup> (Meng et al., 2017)	21.7	17.4	23.7	33.1	21	41.3	32.7	27.8	27.3
BLING-KPE <sup>†</sup> (Xiong et al., 2019)	28.5	22.0	30.3	39.0	27.0	48.1	-	-	-
SKE-Base-Cls <sup>†</sup> (Mu et al., 2020)	-	-	-	-	-	-	39.2	33.0	-
BERT-Span* (Sun et al., 2021)	34.1	28.9	34.0	49.2	29.3	59.3	39.3	32.5	38.3
BERT-SeqTag* (Sun et al., 2021)	37.0	31.5	37.4	54.1	31.8	64.2	40.7	33.5	41.2
ChunkKPE* (Sun et al., 2021)	37.0	31.4	37.0	53.3	31.1	62.7	41.2	33.7	40.9
RankKPE* (Sun et al., 2021)	36.9	31.5	38.1	55.1	32.5	65.5	41.3	34.0	41.8
JointKPE* (Sun et al., 2021)	37.2	31.8	38.2	55.2	32.6	<b>65.7</b>	41.1	33.8	<u>41.9</u>
JointKPE <sup>†</sup> (Sun et al., 2021)	37.1	31.5	38.4	55.5	32.6	<b>65.7</b>	41.1	33.8	<u>41.9</u>
KIEMP <sup>†</sup> (Song et al., 2021)	36.9	29.8	<b>39.2</b>	51.7	<b>34.0</b>	61.5	<b>42.1</b>	<b>34.5</b>	41.2
HyperMatch <sup>†</sup> (Song et al., 2022)	36.4	29.5	39.0	51.5	33.7	61.2	41.6	34.3	40.9
<b>LLM (zero shot)</b>									
ChatGLM2-6b <sup>†</sup> (Song et al., 2023)	16.0	-	11.0	-	8.6	-	-	-	-
GPT-3.5-turbo* (OpenAI, 2022)	20.8	17.0	20.4	26.9	16.6	30.0	13.5	10.8	19.5
<b>Diff-KPE</b>	<b>37.8</b>	<b>32.2</b>	38.5	<b>55.6</b>	32.7	<b>65.7</b>	41.7	34.3	<b>42.3</b>

Table 2: Overall performance of extractive KPE models on OpenKP development set and KP20k testing set. **Bold** indicates the best results, and underlined are the SOTA baselines. <sup>†</sup> indicates results are copied from corresponding papers, and \* are from our reproduction. Note that HyperMatch and KIEMP use the RoBERTa (Liu et al., 2019) as backbone, while others use BERT-base (Devlin et al., 2018).

Model	SemEval-2010		SemEval-2017		Nus		Inspec		Krapivin		Average
	F1@5	F1@10	F1@5	F1@10	F1@5	F1@10	F1@5	F1@10	F1@5	F1@10	
TF-IDF	12.0	18.4	-	-	13.9	18.1	22.3	30.4	11.3	14.3	-
TextRank	17.2	18.1	-	-	19.5	19.0	22.9	27.5	17.2	14.7	-
JointKPE*	28.2	31.0	29.6	<b>37.7</b>	33.9	35.0	31.8	35.0	33.3	29.2	32.4
ChatGLM2-6b <sup>†</sup>	13.2	13.8	-	-	-	-	25.1	30.1	-	-	-
GPT-3.5-turbo*	11.2	11.1	17.0	25.8	14.8	15.2	30.6	33.9	12.8	11.7	18.4
<b>Diff-KPE</b>	<b>29.3</b>	31.0	<b>29.7</b>	37.2	<b>35.2</b>	<b>36.0</b>	<b>32.3</b>	35.0	<b>35.0</b>	<b>31.4</b>	<b>33.2</b>

Table 3: Evaluation results on five small scientific testing sets. The results are evaluated using the models trained on KP20k. Bold indicates the best results. \* results are obtained from our reproduction.

## 5.1. Overall Performance

Table 2 shows the evaluation results of Diff-KPE and baselines. Based on the results, it is evident that the neural KPE methods outperform all the traditional KPE algorithms. Among the traditional methods, the unsupervised methods TF-IDF and TextRank show stable performance on both OpenKP and KP20k datasets, while the feature-based methods PROD and Maui outperform them on OpenKP and KP20k respectively. This is not surprising, as supervised methods benefit from large annotated data during training.

For neural KPE methods, CopyRNN performs the worst as it also focuses on generating abstract keyphrases. HyperMatch, JointKPE and its variant RankKPE show powerful performance, outperforming other baselines such as the phrase classification-based models BLING-KPE, SKE-

Base-Cls, BERT-Span, and the sequence tagging method BERT-SeqTag. It is worth noting that BERT-SeqTag and ChunkKPE exhibit competitive performance compared to RankKPE, indicating their robustness and strong performance.

Overall, Diff-KPE outperforms all baselines excluding KIEMP on both OpenKP and KP20K datasets. Compared to JointKPE, Diff-KPE shows slight improvements in F1@3 and F1@5 but a dramatic improvement in F1@1. Compared to the previous SOTA neural baseline method KIEMP, KIEMP outperforms our Diff-KPE in most F1@k scores on OpenKP and KP20k. However, Diff-KPE still exhibits performance improvements in F1@1, R@1, R@3, and R@5 on OpenKP. We hypothesize that the improvements in Recall benefit from our diffusion module, which is able to inject generated keyphrase embeddings into phrase representations, thereby enhancing the recall performance.

Setting	F1@1	P@1	R@1	F1@3	P@3	R@3	F1@5	P@5	R@5
<b>Diff-KPE</b>	<b>37.8</b>	<b>51.4</b>	<b>32.2</b>	<b>38.5</b>	<b>31.4</b>	<b>55.6</b>	<b>32.7</b>	<b>22.8</b>	<b>65.7</b>
- w/o VIB	36.5	49.4	31.09	37.7	30.8	54.5	32.1	22.4	64.8
- w/o diffusion	36.6	49.7	31.2	37.9	31.0	54.8	32.3	22.5	65.0

Table 4: Evaluation metrics on the OpenKP development set by different settings. “w/o VIB” means Diff-KPE without VIB module, “w/o diffusion” means Diff-KPE without diffusion module.

Moreover, to verify the robustness of Diff-KPE, we also evaluate our model trained with the KP20k dataset on five additional small scientific datasets, as shown in Table 3<sup>3</sup>. Diff-KPE demonstrates better or competitive results on all datasets compared to the best baseline JointKPE. We believe this phenomenon arises from the benefit of the diffusion module: during inference, the diffusion model can generate candidate keyphrase embeddings, providing keyphrase information for the ranking network to better rank each phrase.

## 5.2. Ablation Study

To understand the effect of each component on our Diff-KPE model. We perform the ablation study on the OpenKP development set as following settings:

- - w/o VIB: replace the VIB model with a single feedforward layer for keyphrase classification.
- - w/o diffusion: the diffusion model is removed, and only use the phrase representations obtained from CNNs for ranking and classifying.
- Diff-KPE: the original full joint model.

As shown in Table 4, the absence of the diffusion model or VIB model results in a dramatic drop in performance across all metrics, particularly in F1@1 (1.2 and 1.3 respectively). This performance decline indicates the crucial role of both the diffusion and VIB models in keyphrase ranking. The strong performance of Diff-KPE can be attributed to two main advantages. Firstly, the diffusion module directly incorporates the semantic information of keyphrases into the final phrase representations. Secondly, the supervised VIB module introduces an external classification loss during training, which indirectly enhances the diffusion module or CNNs to generate more informative n-gram embeddings. Therefore, it is evident that the addition of the diffusion module and supervised VIB greatly contributes to the overall performance improvement.

## 6. Case Study

To further demonstrate the effectiveness of the diffusion module in Diff-KPE, we provide examples of

<sup>3</sup>We cannot evaluate KIEMP due to lack of open source code and models (Song et al., 2021).

the extracted keyphrases from our different models (Diff-KPE and Diff-KPE without diffusion module). Two typical cases from the development set of OpenKP are shown in Table 5.

In case (1), both Diff-KPE and Diff-KPE without diffusion successfully extract the desired reference keyphrases “adventure time” and “ryan north” within their top 5 ranked prediction phrases. However, Diff-KPE ranks the phrase “ryan north” higher, resulting in a higher F1@3 score in this case. This illustrates that adding the diffusion module helps the desired keyphrase representation obtain a higher rank score.

Similarly, in case (2), Diff-KPE ranks the desired keyphrases “codesnip” and “oracle script” higher compared to the model without diffusion. As a result, Diff-KPE successfully extracts all the reference keyphrases in case (2). The main reason for these results may be that the keyphrase embeddings generated by the diffusion module are directly injected into each phrase representation, enabling the ranking network to better rank each phrase by utilizing the keyphrase information.




Figure 2: T-SNE visualization of phrase embeddings from OpenKP dataset.

We also analyze the generated keyphrase embeddings quality. We apply T-SNE (Van der Maaten and Hinton, 2008) to reduce all the phrase representation’s dimensions to 2 in Figure 2. We can find that the oracle keyphrases (green dots) and generated keyphrases (blue dots) are clustered together and far away from most non-keyphrase



<p><b>(1) Partial Document:</b>  ... in Comics RealWorld Objects Non canon <b>Adventure Time</b> comic English Share Adventure Time is a comic book series published by BOOM Studios written by Dinosaur Comics creator <b>Ryan North</b>, and illustrated by Shelli Paroline and Braden Lamb. The comic book is released monthly beginning with issue 1 in February 2012 ... (URL: <a href="http://adventuretime.wikia.com/wiki/Adventure_Time_(comic)">http://adventuretime.wikia.com/wiki/Adventure_Time_(comic)</a>)</p> <p><b>Reference Keyphrases:</b>  <b>adventure time</b>; <b>ryan north</b></p>
<p><b>Without diffusion module:</b>  <b>adventure time</b>; boom studios; comic book series; dinosaur comics; <b>ryan north</b></p> <p><b>Diff-KPE:</b>  <b>adventure time</b>; comic book series; <b>ryan north</b>; comic book; dinosaur comics</p>
<p><b>(2) Partial Document:</b>  <b>CodeSnip:</b> How to Run Any <b>Oracle Script</b> File Through Shell Script in UNIX ... by Deepankar Sarangi ... Listing 1 ... The first line is a comment line which is UNIX kernel specific. In the following approach the available shell is KORN shell ... (URL: <a href="http://aspalliance.com/1589_CodeSnip_How_to_Run_Any_Oracle_Script_File_Through_Shell_Script_in_UNIX.4">http://aspalliance.com/1589_CodeSnip_How_to_Run_Any_Oracle_Script_File_Through_Shell_Script_in_UNIX.4</a>) <b>Reference Keyphrases:</b>  <b>codesnip</b>; <b>oracle script</b></p>
<p><b>Without diffusion module:</b>  shell script; oracle script file through; oracle script file; shell scripts; <b>codesnip</b></p> <p><b>Diff-KPE:</b>  unix; shell script; <b>codesnip</b>; oracle script file through; <b>oracle script</b></p>

Table 5: Example of keyphrase extraction results on two selected OpenKP development examples. The phrase in red is the desired reference keyphrase.

embeddings (red dots). This finding demonstrates that our diffusion model is powerful in recovering keyphrase embeddings.

## 7. Conclusion

In this paper, we propose Diff-KPE, a novel joint keyphrase extraction (KPE) model composed of three essential modules: the diffusion module, the ranking network, and a supervised VIB module. Each component plays a crucial role in learning expressive phrase representations. The diffusion module is responsible for generating keyphrase embeddings, effectively infusing keyphrase semantic information into the final phrase representation. Simultaneously, the supervised VIB introduces a classification loss for each phrase, encouraging the model to generate more informative representations and ultimately improving the ranking performance. Experimental results on seven keyphrase extraction benchmark datasets demonstrate the effectiveness and superiority of Diff-KPE.

However, since our model requires many steps of forward noise injection and backward denoising, our Diff-KPE is about 2x slower than the previous SOTA model JointKPE during inference. Moreover, our model also lacks the ability to generate abstractive keyphrases. In future work, we plan to improve the computation efficiency and explore the application of Diff-KPE in abstractive keyphrase generation, leveraging its powerful architecture and flexibility for generating concise and informative keyphrases.

## 8. Ethics Statement

We take ethical considerations seriously and strictly adhere to the Ethics Policy. This paper focuses on the attempt to the application of diffusion model for keyphrase extraction. Both the datasets and base models used in this paper are publicly available and have been widely adopted by researchers. We ensure that the findings and conclusions of this paper are reported accurately and objectively.

Rabah Alzaidy, Cornelia Caragea, and C Lee Giles. 2019. Bi-lstm-crf sequence labeling for keyphrase extraction from scholarly documents. In *The world wide web conference*, pages 2551–2557.

Isabelle Augenstein, Mrinal Das, Sebastian Riedel, Lakshmi Vikraman, and Andrew McCallum. 2017. Semeval 2017 task 10: Scienceie-extracting keyphrases and relations from scientific publications. *arXiv preprint arXiv:1704.02853*.

Florian Boudin. 2018. Unsupervised keyphrase extraction with multipartite graphs. *arXiv preprint arXiv:1803.08721*.

Ricardo Campos, Vítor Mangaravite, Arian Pasquali, Alípio Mário Jorge, Célia Nunes, and Adam Jatowt. 2018. A text feature based automatic keyword extraction method for single documents. In *European conference on information retrieval*, pages 684–691. Springer.

- Jun Chen, Xiaoming Zhang, Yu Wu, Zhao Yan, and Zhoujun Li. 2018. Keyphrase generation with correlation constraints. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4057–4066.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Haoran Ding and Xiao Luo. 2021. Attentionrank: Unsupervised keyphrase extraction using self and cross attentions. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1919–1928.
- Samhaa R El-Beltagy and Ahmed Rafea. 2009. Kp-miner: A keyphrase extraction system for english and arabic documents. *Information systems*, 34(1):132–144.
- Corina Florescu and Cornelia Caragea. 2017a. A new scheme for scoring phrases in unsupervised keyphrase extraction. In *Advances in Information Retrieval: 39th European Conference on IR Research, ECIR 2017, Aberdeen, UK, April 8-13, 2017, Proceedings 39*, pages 477–483. Springer.
- 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.
- Shansan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and LingPeng Kong. 2022. Diffuseq: Sequence to sequence text generation with diffusion models. *arXiv preprint arXiv:2210.08933*.
- Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. 2022. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*.
- Danqing Huang, Jing Liu, Chin-Yew Lin, and Jian Yin. 2018a. [Neural math word problem solver with reinforcement learning](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 213–223, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Danqing Huang, Shuming Shi, Chin-Yew Lin, and Jian Yin. 2017. [Learning fine-grained expressions to solve math word problems](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 805–814, Copenhagen, Denmark. Association for Computational Linguistics.
- Danqing Huang, Jin-Ge Yao, Chin-Yew Lin, Qingyu Zhou, and Jian Yin. 2018b. [Using intermediate representations to solve math word problems](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 419–428, Melbourne, Australia. 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.
- Su Nam Kim, Olena Medelyan, Min-Yen Kan, and Timothy Baldwin. 2013. Automatic keyphrase extraction from scientific articles. *Language resources and evaluation*, 47:723–742.
- Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. 2020. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*.
- Mikalai Krapivin, Aliaksandr Autsaev, and Maurizio Marchese. 2009. Large dataset for keyphrases extraction.
- Mayank Kulkarni, Debanjan Mahata, Ravneet Arora, and Rajarshi Bhowmik. 2022. Learning rich representation of keyphrases from text. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 891–906.
- Da-Wei Li, Danqing Huang, Tingting Ma, and Chin-Yew Lin. 2021. Towards topic-aware slide generation for academic papers with unsupervised mutual learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13243–13251.
- Haoran Li, Junnan Zhu, Jiajun Zhang, Chengqing Zong, and Xiaodong He. 2020. Keywords-guided abstractive sentence summarization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 8196–8203.
- Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. 2022. Diffusion-lm improves controllable text generation. *Advances in Neural Information Processing Systems*, 35:4328–4343.

- Xiang Lisa Li and Jason Eisner. 2019. Specializing word embeddings (for parsing) by information bottleneck. *arXiv preprint arXiv:1910.00163*.
- Xinnian Liang, Shuangzhi Wu, Mu Li, and Zhoujun Li. 2021. Unsupervised keyphrase extraction by jointly modeling local and global context. *arXiv preprint arXiv:2109.07293*.
- 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. *arXiv preprint arXiv:1907.11692*.
- Yi Luan, Mari Ostendorf, and Hannaneh Hajishirzi. 2017. Scientific information extraction with semi-supervised neural tagging. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2641–2651.
- Debanjan Mahata, John Kuriakose, Rajiv Shah, and Roger Zimmermann. 2018. Key2vec: Automatic ranked keyphrase extraction from scientific articles using phrase embeddings. 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 634–639.
- Roberto Martínez-Cruz, Alvaro J López-López, and José Portela. 2023. Chatgpt vs state-of-the-art models: A benchmarking study in keyphrase generation task. *arXiv preprint arXiv:2304.14177*.
- Olena Medelyan, Eibe Frank, and Ian H Witten. 2009. Human-competitive tagging using automatic keyphrase extraction. Association for Computational Linguistics.
- Rui Meng, Sanqiang Zhao, Shuguang Han, Daqing He, Peter Brusilovsky, and Yu Chi. 2017. Deep keyphrase generation. *arXiv preprint arXiv:1704.06879*.
- 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.
- 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. *arXiv preprint arXiv:2002.05407*.
- Nhung TH Nguyen, Makoto Miwa, and Sophia Ananiadou. 2023. Span-based named entity recognition by generating and compressing information. *arXiv preprint arXiv:2302.05392*.
- Thuy Dung Nguyen and Min-Yen Kan. 2007. Keyphrase extraction in scientific publications. In *International conference on Asian digital libraries*, pages 317–326. Springer.
- OpenAI. 2022. [ChatGPT](#).
- Aitor Ormazabal, Mikel Artetxe, Aitor Soroa, Gorka Labaka, and Eneko Agirre. 2022. Principled paraphrase generation with parallel corpora. *arXiv preprint arXiv:2205.12213*.
- Martin F Porter. 1980. An algorithm for suffix stripping. *Program*, 14(3):130–137.
- Nils Reimers and Iryna Gurevych. 2019. SentenceBERT: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695.
- Dhruva Sahrawat, Debanjan Mahata, Haimin Zhang, Mayank Kulkarni, Agniv Sharma, Rakesh Gosangi, Amanda Stent, Yaman Kumar, Rajiv Ratn Shah, and Roger Zimmermann. 2020. Keyphrase extraction as sequence labeling using contextualized embeddings. In *Advances in Information Retrieval: 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14–17, 2020, Proceedings, Part II 42*, pages 328–335. Springer.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR.
- Mingyang Song, Yi Feng, and Liping Jing. 2022. Hyperbolic relevance matching for neural keyphrase extraction. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5710–5720.
- Mingyang Song, Xuelian Geng, Songfang Yao, Shilong Lu, Yi Feng, and Liping Jing. 2023. Large language models as zero-shot keyphrase extractor: A preliminary empirical study. *arXiv preprint arXiv:2312.15156*.
- Mingyang Song, Liping Jing, and Lin Xiao. 2021. Importance estimation from multiple perspectives for keyphrase extraction. *arXiv preprint arXiv:2110.09749*.

- Karen Sparck Jones. 1972. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*, 28(1):11–21.
- Si Sun, Zhenghao Liu, Chenyan Xiong, Zhiyuan Liu, and Jie Bao. 2021. Capturing global informativeness in open domain keyphrase extraction. In *Natural Language Processing and Chinese Computing: 10th CCF International Conference, NLPCC 2021, Qingdao, China, October 13–17, 2021, Proceedings, Part II 10*, pages 275–287. Springer.
- Yi Sun, Hangping Qiu, Yu Zheng, Zhongwei Wang, and Chaoran Zhang. 2020. Sifrank: a new baseline for unsupervised keyphrase extraction based on pre-trained language model. *IEEE Access*, 8:10896–10906.
- Naftali Tishby, Fernando C Pereira, and William Bialek. 2000. The information bottleneck method. *arXiv preprint physics/0004057*.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(11).
- Slava Voloshynovskiy, Mouad Kondah, Shideh Rezaeifar, Olga Taran, Taras Holotyak, and Danilo Jimenez Rezende. 2019. Information bottleneck through variational glasses. *arXiv preprint arXiv:1912.00830*.
- Xiao Wang, Shihan Dou, Limao Xiong, Yicheng Zou, Qi Zhang, Tao Gui, Liang Qiao, Zhazhan Cheng, and Xuanjing Huang. 2022. Miner: Improving out-of-vocabulary named entity recognition from an information theoretic perspective. *arXiv preprint arXiv:2204.04391*.
- Yansen Wang, Zhen Fan, and Carolyn Rose. 2020. Incorporating multimodal information in open-domain web keyphrase extraction. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1790–1800.
- Peter West, Ari Holtzman, Jan Buys, and Yejin Choi. 2019. Bottlesum: Unsupervised and self-supervised sentence summarization using the information bottleneck principle. *arXiv preprint arXiv:1909.07405*.
- Lee Xiong, Chuan Hu, Chenyan Xiong, Daniel Campos, and Arnold Overwijk. 2019. Open domain web keyphrase extraction beyond language modeling. *arXiv preprint arXiv:1911.02671*.
- Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Yingxia Shao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. 2022. Diffusion models: A comprehensive survey of methods and applications. *arXiv preprint arXiv:2209.00796*.
- Hongyi Yuan, Zheng Yuan, Chuanqi Tan, Fei Huang, and Songfang Huang. 2022. Seqdiffuseq: Text diffusion with encoder-decoder transformers. *arXiv preprint arXiv:2212.10325*.
- Xingdi Yuan, Tong Wang, Rui Meng, Khushboo Thaker, Peter Brusilovsky, Daqing He, and Adam Trischler. 2018. One size does not fit all: Generating and evaluating variable number of keyphrases. *arXiv preprint arXiv:1810.05241*.
- Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. 2022. Glm-130b: An open bilingual pre-trained model. *arXiv preprint arXiv:2210.02414*.
- Cenyuan Zhang, Xiang Zhou, Yixin Wan, Xiaoqing Zheng, Kai-Wei Chang, and Cho-Jui Hsieh. 2022. Improving the adversarial robustness of nlp models by information bottleneck. *arXiv preprint arXiv:2206.05511*.
- Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023. Diffusum: Generation enhanced extractive summarization with diffusion. *arXiv preprint arXiv:2305.01735*.
- Linhan Zhang, Qian Chen, Wen Wang, Chong Deng, Shiliang Zhang, Bing Li, Wei Wang, and Xin Cao. 2021. Mderank: A masked document embedding rank approach for unsupervised keyphrase extraction. *arXiv preprint arXiv:2110.06651*.
- Qi Zhang, Yang Wang, Yeyun Gong, and Xuan-Jing Huang. 2016. Keyphrase extraction using deep recurrent neural networks on twitter. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pages 836–845.
- Qingyu Zhou and Danqing Huang. 2019. [Towards generating math word problems from equations and topics](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 494–503, Tokyo, Japan. Association for Computational Linguistics.