

WR-ONE2SET: Towards Well-Calibrated Keyphrase Generation

Binbin Xie^{1,3}, Xiangpeng Wei², Baosong Yang², Huan Lin², Jun Xie²,
Xiaoli Wang³, Min Zhang⁴ and Jinsong Su^{1,3*}

¹School of Informatics, Xiamen University, China ²Alibaba Group, China

³Key Laboratory of Digital Protection and Intelligent Processing of Intangible Cultural Heritage of Fujian and Taiwan, Ministry of Culture and Tourism, China

⁴Soochow University, China

xdbl@stu.xmu.edu.cn pemywei@gmail.com

yangbaosong.ybs@alibaba-inc.com sjs@xmu.edu.cn

Abstract

Keyphrase generation aims to automatically generate short phrases summarizing an input document. The recently emerged ONE2SET paradigm (Ye et al., 2021) generates keyphrases as a set and has achieved competitive performance. Nevertheless, we observe serious calibration errors outputted by ONE2SET, especially in the *over-estimation* of \emptyset token (means “no corresponding keyphrase”). In this paper, we deeply analyze this limitation and identify two main reasons behind: 1) the parallel generation has to introduce excessive \emptyset as padding tokens into training instances; and 2) the training mechanism assigning target to each slot is unstable and further aggravates the \emptyset token over-estimation. To make the model well-calibrated, we propose WR-ONE2SET which extends ONE2SET with an adaptive instance-level cost **W**eighting strategy and a target **R**e-assignment mechanism. The former dynamically penalizes the over-estimated slots for different instances thus smoothing the uneven training distribution. The latter refines the original inappropriate assignment and reduces the supervisory signals of over-estimated slots. Experimental results on commonly-used datasets demonstrate the effectiveness and generality of our proposed paradigm.

1 Introduction

Keyphrases are short phrases fully encoding the main information of a given document. They can not only facilitate readers to quickly understand the document, but also provide useful information to many downstream tasks, including document classification (Hulth and Megyesi, 2006), summarization (Wang and Cardie, 2013), etc.

With the rapid development of deep learning, *keyphrase generation* (Meng et al., 2017) has attracted increasing attention due to its ability to produce phrases that even do not match any contiguous

* Corresponding author.

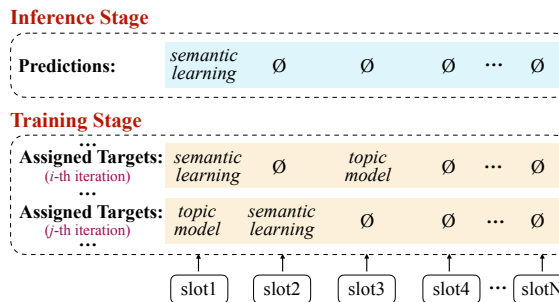


Figure 1: An example of ONE2SET paradigm at training and inference stages. “Assigned Targets ($*$ -th iteration)” represents the multiple feasible target permutations generated by K -step target assignment mechanism at different training iterations. In this case, both “slot2” and “slot3” are expected to generate keyphrases. However, they often use \emptyset token as supervisory signals, and thus over-estimate and output \emptyset token.

subsequence of the source document.¹ Dominant models of keyphrase generation

are constructed under three paradigms: ONE2ONE (Meng et al., 2017), ONE2SEQ (Yuan et al., 2020) and ONE2SET (Ye et al., 2021). Among these paradigms, ONE2SET exhibits the state-of-the-art (SOTA) performance. As illustrated in Figure 1, it considers keyphrase generation as a set generation task. After padding keyphrases to a fixed number with special token \emptyset , they define multiple slots that individually generate each keyphrase in parallel. During training, each slot is assigned with a keyphrase or \emptyset token² via a K -step target assignment mechanism. Specifically, the model first generates K tokens from each slot and then determines the optimal target assignment using a bipartite matching algorithm (Kuhn, 2010). The superiority of ONE2SET stems from its conditional independence, that is, the prediction distribution of each slot depends only on the given document other than the order of keyphrases like

¹An example is shown in Appendix, Table A.1.

²In this work, we define that the keyphrase can not be a \emptyset token.

ONE2SEQ. This is more compatible with the unordered property of keyphrases and decreases the difficulty of the model training (Ye et al., 2021).

Despite of its success, we observe serious over-estimation problem on \emptyset token, which significantly affects the generation quality. For example, in Figure 1, both “slot2” and “slot3” are expected to generate keyphrases, but \emptyset token is over-confidently given. Two questions naturally arise: 1) *what are reasons behind the over-estimation problem in ONE2SET?* and 2) *how can we alleviate them?*

In order to answer the first question, we conduct extensive analyses, and conclude two reasons. Firstly, the over-estimation is a by-product inherently carried by the parallel generation. More concretely, excessive \emptyset tokens have been introduced as the padding tokens and served as supervisory signals in training data. The unbalanced data and the lack of dependency among slots leads each slot to learn to commonly generate \emptyset token. Secondly, the K -step target assignment mechanism provides multiple feasible target permutations that are assigned to slots. As shown in Figure 1, the targets of the given document can be assigned in different permutation at each training iteration, which further increases the probability of \emptyset token to be assigned as supervisory signal for each slot, thus exacerbating the over-estimation problem. Both problems make the learned probabilities of the assigned targets deviate from its ground truth likelihood, finally constructing a miscalibrated model.

Consequently, we approach the above problems from the calibration perspective and propose two strategies that extend ONE2SET to WR-ONE2SET. Specifically, an *adaptive instance-level cost weighting* is first introduced to penalize the over-estimated slots of different instances. According to the seriousness of the issue, instances are rendered different weights, therefore dynamically balancing the model training. Besides, we propose a *target re-assignment* mechanism to refine the original inappropriate assignment and reduce the supervisory signals of \emptyset token. In particular, we re-assign targets for the slots potentially generating fresh keyphrases but being pre-assigned with \emptyset token. In these ways, WR-ONE2SET is encouraged to produce well-calibrated probabilities on keyphrase generation. Overall, major contributions of our work are three-fold:

- Through in-depth analyses, we point out that

the advanced keyphrase generation architecture ONE2SET suffers from the \emptyset token over-estimation, which is inherently caused by its parallelism and the target assignment mechanism.

- We propose WR-ONE2SET which enhances the original framework with two effective strategies to calibrate the over-estimation problem from the training perspective.
- Extensive experiments on five widely-used datasets reveal the universal-effectiveness of our model.
- We release our code at <https://github.com/DeepLearnXMU/WR-One2Set>.

2 Related Work

Early studies mainly focus on automatic keyphrase extraction (Hulth, 2003; Mihalcea and Tarau, 2004; Nguyen and Kan, 2007; Wan and Xiao, 2008), which aims to directly extract keyphrases from the input document. Recently, with the rapid development of deep learning, neural network-based models have been widely used in keyphrase generation. Typically, these models are based on an attentional encoder-decoder framework equipped with copy mechanism, which is able to generate both present and absent keyphrases (Meng et al., 2017). Generally, these models are constructed under the following paradigms: 1) ONE2ONE (Meng et al., 2017; Chen et al., 2019a,b). Under this paradigm, the input document is paired with each target keyphrase to form an independent training instance for model training. During inference, the models are encouraged to produce multiple keyphrases via beam search. 2) ONE2SEQ (Chan et al., 2019; Yuan et al., 2020; Chen et al., 2020; Wu et al., 2021). It considers keyphrase generation as a sequence generation task, where different keyphrases are concatenated into a sequence in a predefined order. In this way, the semantic dependence between keyphrases can be exploited to benefit keyphrase generation. 3) ONE2SET (Ye et al., 2021). Unlike ONE2SEQ, this paradigm considers the keyphrases as a set, which can be predicted from slots in a parallel manner and partial target matching algorithm.

Considering that ONE2ONE neglects the correlation among keyphrases, the most popular paradigm ONE2SEQ exploits the correlation by pre-defining the keyphrases order for model training and inference. Nevertheless, ONE2SEQ is the opposite of the flexible and unordered properties of keyphrases,

increasing the difficulty of the model training. Due to the parallelism and the conditional independence, ONE2SET attracts much attention in the keyphrase generation community, and achieves the SOTA performance. As this method has just been put forward, it is inevitable to exist imperfections. Hence, we are committed to analyses and further optimizing this framework. To the best of our knowledge, this is the first attempt to improve ONE2SET.

3 Background

Here, we briefly introduce SETTRANS (Ye et al., 2021), which is based on the ONE2SET paradigm. It is a Transformer-based, semi-autoregressive model. Typically, it introduces N slots, each of which introduces a learnable control code as the additional decoder input, to generate keyphrases or \emptyset tokens in parallel. Its training involves two stages: 1) a K -step target assignment mechanism is firstly used to determine the correspondence between each prediction and target, and then 2) a new training objective is introduced to optimize the whole model. It contains two set losses to separately deal with two kinds of keyphrases: *present keyphrases* appearing in the input document, and *absent keyphrases* that do not match any contiguous subsequence of the document.

K -Step Target Assignment At this stage, the model predicts K tokens from each slot, where the predicted probability distributions are also collected. Then, an optimal assignment m between predictions and targets can be found by a bipartite matching algorithm (Kuhn, 2010):

$$m = \arg \min_{m \in M(N)} \sum_{i=1}^N \mathcal{C}_{\text{match}}(y^{m(i)}, \mathbf{P}^i), \quad (1)$$

where $M(N)$ denotes a set of all N -length target index permutations, and the optimal permutation m can be considered as a mapping function from the slot i to the target index $m(i)$.³ $\mathcal{C}_{\text{match}}(y^{m(i)}, \mathbf{P}^i)$ is a pair-wise matching loss between the target $y^{m(i)}$ and the predicted probability distributions \mathbf{P}^i of the slot i . Note that the set of targets are also padded to size N with \emptyset tokens.

Model Optimization with Set Losses During the second stage, the model is trained with the sum of two set losses. Concretely, slots are equally

³Please note that instead of following Ye et al. (2021) to use the function $\pi(i')$ mapping the target index i' to the slot index $\pi(i')$, we use $m(i)$ that is the inverse function of $\pi(i)$, so as to facilitate subsequent descriptions.

split into two sets, dealing with the generations of present and absent keyphrases, respectively. Next, the above target assignment is performed on these two sets separately, forming a mapping m^p for present keyphrases, and a mapping m^a for absent keyphrases. Finally, the training objective becomes

$$\mathcal{L}(\theta) = - \left[\sum_{i=1}^{\frac{N}{2}} \mathcal{L}^p(\theta, y^{m^p(i)}) + \sum_{i=\frac{N}{2}+1}^N \mathcal{L}^a(\theta, y^{m^a(i)}) \right] \quad (2)$$

$$\mathcal{L}^p(\theta, z) = \begin{cases} \lambda_{pre} \cdot \sum_{t=1}^{|z|} \log \hat{p}_t^i(z_t) & \text{if } z=\emptyset \\ \sum_{t=1}^{|z|} \log \hat{p}_t^i(z_t) & \text{otherwise} \end{cases} \quad (3)$$

where λ_{pre} is a hyper-parameter used to reduce the negative effect of excessive \emptyset tokens, z_t symbolizes the t -th token of the target z , and \hat{p}_t^i is the t -th predicted probability distribution of the i -th slot using teacher forcing. Meanwhile, $\mathcal{L}^a(\theta, z)$ is defined in the similar way as $\mathcal{L}^p(\theta, z)$ with a hyper-parameter λ_{abs} .

4 Preliminary Analyses

Although ONE2SET has achieved competitive performance, it still faces one major problem, i.e. *\emptyset token over-estimation*. This occurs in such slots that produce \emptyset tokens via the vanilla prediction while are able to generate correct keyphrases through the non- \emptyset prediction⁴. For illustration, we force all slots to generate non- \emptyset predictions during inference, where 14.6% of slots can produce correct ones. However, if we remove this restriction, 34.5% of these slots directly output \emptyset tokens, revealing the over-estimation of \emptyset token. Such kind of miscalibration (Guo et al., 2017; Kumar and Sarawagi, 2019) is a common drawback in neural network based models, which not only seriously hurts the generation quality of the ONE2SET paradigm, but also limits the users' trust towards it.

To understand the reasons behind this, we use the commonly-used KP20k dataset (Meng et al., 2017) to train a standard SETTRANS model, where the assigned targets to the slots of each instance are recorded during the last 80,000 training steps with an interval of 8,000 steps. Here, we can obtain two crucial observations.

Observation 1: *Excessive \emptyset tokens have been introduced as the padding tokens and served as supervisory signals in training data.* ONE2SET models keyphrase generation in a parallel computation fashion, therefore extensive padding \emptyset to

⁴When performing the non- \emptyset prediction, we remove \emptyset token from the prediction vocabulary to generate a keyphrase.

Target Type	Pre.KP Slots	Abs.KP Slots
\emptyset	72.4%	80.4%
Target KP	27.6%	19.6%

Table 1: The proportions of \emptyset token and target keyphrases used as supervisory signals during training. ‘‘KP’’ means keyphrase. ‘‘Pre.KP’’ and ‘‘Abs.KP’’ represent present and absent keyphrases, respectively.

Instance Type	All KP Slots
Instance(#OV-Slot=0)	42.1%
Instance(#OV-Slot=1)	31.9%
Instance(#OV-Slot=2)	15.5%
Instance(#OV-Slot \geq 3)	10.5%

Table 2: The proportions of instances involving different numbers of slots over-estimating \emptyset token. Instance(#OV-Slot= n) means the instances containing n slots over-estimating \emptyset token. Please note that the greater n , the more severe \emptyset token over-estimation.

kens are used to make sure the fixed lengths of different samples. Table 1 shows the proportions of \emptyset token and target keyphrases involved during the model training. We can observe that on both present and absent keyphrase slots, \emptyset token accounts for the vast majority, exceeding 70%. In addition, instances suffer from different degrees of \emptyset token over-estimation. Table 2 shows the proportions of training instances grouped by the number of slots over-estimating \emptyset token. We can find that the instances (e.g. Instance(#OV-Slot \geq 1)) account for significant proportions, and exist varying degrees of \emptyset token over-estimation.

Observation 2: *The K -step assignment mechanism is unstable and further increases the possibility of \emptyset tokens being served as supervisory signals for some slots.* In spite of the clever design of K -step assignment mechanism, it unstably provides different feasible target permutations to slots at the training time. We argue that this further widens the gap between the distribution of supervisory signals and that of the ground-truth.

To illustrate this, we classify the slots of each instance into three categories according to its target assignments: 1) Slot(\emptyset), each slot of this category is always assigned with \emptyset tokens. Apparently, these slots hardly generate keyphrases after training; 2) Slot(Target KP), each slot of this category is always assigned with target keyphrases and thus it has high probability of generating a keyphrase; 3) Slot(\emptyset +Target KP), each slot is assigned with

Slot Type	Pre.KP Slots	Abs.KP Slots
Slot(\emptyset)	61.2%	66.4%
Slot(Target KP)	17.6%	9.3%
Slot(\emptyset +Target KP)	21.2%	24.4%

Table 3: The proportions of slots with different target assignments during the model training. Slot(\emptyset +Target KP) means the slots are assigned with \emptyset tokens and target keyphrases alternatively at different iterations. Note that the higher proportions of Slot(\emptyset +Target KP), the more slots contain unstable supervisory signals.

target keyphrases or \emptyset tokens at different iterations during model training. From Table 3, we can observe that on both present and absent keyphrase slots, the proportions of Slot(\emptyset +Target KP) are quite high, exceeding those of Slot(Target KP). Quite evidently, the supervisory signals of slots in Slot(\emptyset +Target KP) are unstable. Those slots that should be labeled with Target KP are assigned with \emptyset token, further decreasing the probabilities of these slots generating keyphrases.

5 WR-ONE2SET

As discussed above, the parallelism and the training mechanism of ONE2SET bring the advantages of conditional independence, but inherently lead to the miscalibration of the model. Our principle is to maintain the primary advantages, and meanwhile, calibrating the model with lightweight strategies. To this end, we propose WR-ONE2SET that significantly extends the conventional ONE2SET paradigm in two training aspects, including an *adaptive instance-level cost weighting strategy*, and a *target re-assignment mechanism*.

To facilitate subsequent descriptions, we summarize all related formal definitions in Appendix, Table A.2 for better understanding this paradigm.

5.1 Adaptive Instance-Level Cost Weighting

Connection to Observation 1. As analyzed previously, excessive \emptyset tokens lead to the over-estimation of \emptyset token. Although SETTRANS introduces hyper-parameters λ_{pre} and λ_{abs} to adjust the training loss of conventional ONE2SET paradigm, such fixed hyper-parameters are still unable to deal with this issue well due to the different degrees of \emptyset token over-estimation in different training instances.

We alternatively develop an adaptive instance-level cost weighting strategy to dynamically scale the losses corresponding to \emptyset tokens, alleviating

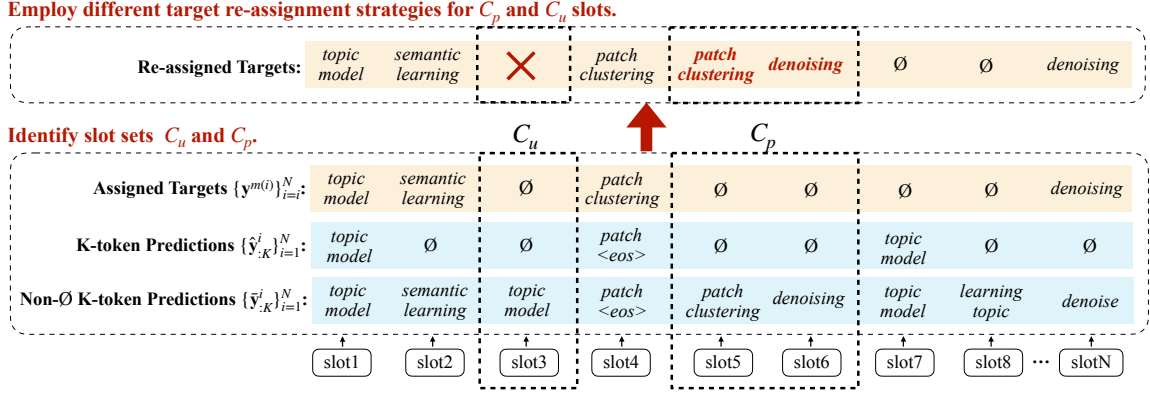


Figure 2: The procedure of target re-assignment involving two steps: identifying the potential slot set C_p and the unimportant one C_u , and then employing different re-assignment operations to deal with them, respectively. Here, “✗” represents assigning no supervisory signal to the slots of C_u , and following Ye et al. (2021), we set $K = 2$.

the class imbalance of training data. Concretely, we first identify a set of slots, denoted as $C_{1\emptyset}$, where each slot is assigned with a keyphrase as supervisory signal. Intuitively, for each slot i in $C_{1\emptyset}$, the degree of \emptyset token over-estimation is related to its two predicted probabilities using teacher forcing (See Section 4): 1) $\hat{p}^i(y_0^{m(i)})$, symbolizing the predicted probability of the first token of assigned target, and 2) $\hat{p}^i(\emptyset)$, denoting the predicted probability of \emptyset token. Thus, we directly use the ratio between $\hat{p}^i(y_0^{m(i)})$ and $\hat{p}^i(\emptyset)$ to approximately quantify the degree of \emptyset token over-estimation for training efficiency. Furthermore, we define this degree for each instance as

$$\lambda_{adp} = \frac{1}{|C_{1\emptyset}|} \cdot \sum_{i \in C_{1\emptyset}} \min\left(\frac{\hat{p}^i(y_0^{m(i)})}{\hat{p}^i(\emptyset)}, 1\right). \quad (4)$$

Note that for each slot i in $C_{1\emptyset}$, if its predicted probability $\hat{p}^i(y_0^{m(i)})$ is greater than $\hat{p}^i(\emptyset)$, it is considered to have no \emptyset token over-estimation, and we directly limit its ratio to 1.

Finally, we adjust the hyper-parameters λ_{pre} and λ_{abs} of Equation 3 into $\lambda_{adp} \cdot \lambda_{pre}$ and $\lambda_{adp} \cdot \lambda_{abs}$ for each training instance, respectively. Note that, λ_{adp} is dynamically updated during the training process, and thus is more general for model training compared with fixed hyper-parameters.

5.2 Target Re-Assignment

Connection to Observation 2. Due to the effect of K -step target assignment mechanism, many slots are alternatively assigned with target keyphrases and \emptyset tokens, which decreases the probabilities of these slots generating correct keyphrases.

We propose a target re-assignment mechanism to alleviate this issue. As shown in the upper part

of Figure 2, during the process of K -step target assignment, we first record three kinds of phrases for each slot i : 1) $y^{m(i)}$, the assigned target of the slot i ; 2) $\hat{y}_{:K}^i$, the first K tokens of the vanilla prediction from the slot i . Note that $\hat{y}_{:K}^i$ may be a \emptyset token; and 3) $\bar{y}_{:K}^i$, the first K tokens of the non- \emptyset prediction from the slot i .

Here, we mainly focus on the slots, each of which is assigned with ϕ token as supervisory signals and its non- ϕ K -token prediction is consistent with some targets. For such slot i , if its non- \emptyset K -token prediction $\bar{y}_{:K}^i$ is totally different from all K -token predictions $\{\hat{y}_{:K}^i\}_{i=1}^N$, we consider it has the potential to generate a fresh keyphrase and boost the model performance. Thus, we include it into the **potential slot set** C_p . By contrast, if its $\bar{y}_{:K}^i$ has occurred in the set of $\{\hat{y}_{:K}^i\}_{i=1}^N$, we regard it as an unimportant slot without effect on the model performance, and add it into the **unimportant slot set** C_u . Back to Figure 2, we observe that the non- \emptyset K -token prediction of “slot3” is “topic model”, which is also the K -token prediction of “slot1” and “slot7”. Thus, “slot3” is an unimportant slot. Meanwhile, both “slot5” and “slot6” are potential slots.

Then, as illustrated in the lower part of Figure 2, we employ two target re-assignment operations to deal with the above two kinds of slots, respectively: 1) we re-assign each slot of C_p with its best-matched target keyphrase, so as to increase the probability of this slot generating the target keyphrase; and 2) we assign no target to each slot of C_u , which alleviates the problem that the same target is assigned to different slots as supervisory signals. In this way, the training losses of slots in C_u will be masked at this training iteration. Let us

Model	Inspec		NUS		Krapivin		SemEval		KP20k	
	$F_1@5$	$F_1@M$	$F_1@5$	$F_1@M$	$F_1@5$	$F_1@M$	$F_1@5$	$F_1@M$	$F_1@5$	$F_1@M$
<i>Existing Neural Keyphrase Generation Models</i>										
CATSEQ(R)	0.225	0.262	0.323	0.397	0.269	0.354	0.242	0.283	0.291	0.367
CATSEQ	0.281 ₅	0.325 ₆	0.370 ₇	0.419 ₁₀	0.315 ₈	0.365 ₅	0.287 ₁₄	0.325 ₁₅	0.332 ₁	0.377 ₁
UNIKEYPHRASE	0.260	0.288	0.415	0.443	—	—	0.302	0.322	0.347	0.352
SETTRANS	0.285 ₃	0.324 ₃	0.406 ₁₂	0.450 ₇	0.326 ₁₂	0.364 ₁₂	0.331 ₂₀	0.357 ₁₃	0.358 ₅	0.392 ₄
PROMPTKP	0.260	0.294	0.412	0.439	—	—	0.329	0.356	0.351	0.355
<i>Our Models</i>										
SETTRANS	0.282 ₂	0.320 ₂	0.399 ₅	0.437 ₈	0.334 ₇	0.368 ₄	0.333 ₆	0.357 ₄	0.359 ₂	0.392 ₂
SETTRANS(w/o $\lambda_{pre}, \lambda_{abs}$)	0.100 ₄	0.148 ₆	0.173 ₁₈	0.258 ₄₂	0.155 ₉	0.280 ₁₆	0.129 ₇	0.188 ₈	0.191 ₇	0.321 ₉
SETTRANS(#SLOT=12)	0.280 ₁	0.316 ₂	0.387 ₉	0.423 ₄	0.324 ₅	0.369 ₃	0.302 ₆	0.327 ₄	0.347 ₃	0.385 ₃
SETTRANS(#SLOT=16)	0.280 ₇	0.318 ₆	0.384 ₇	0.431 ₁₀	0.319 ₃	0.362 ₁	0.316 ₁₅	0.356 ₁₇	0.342 ₂	0.382 ₂
SETTRANS(#SLOT=24)	0.284 ₁₃	0.320 ₁₈	0.400 ₁	0.453 ₁	0.331 ₁	0.367 ₅	0.327 ₁₈	0.359 ₁₁	0.360 ₁	0.395₁
SETTRANS(#SLOT=28)	0.277 ₉	0.317 ₈	0.402 ₁	0.454₄	0.338 ₄	0.374₆	0.315 ₃	0.349 ₂	0.355 ₄	0.392 ₄
SETTRANS(w/ BATCHING)	0.281 ₇	0.271 ₆	0.379 ₄	0.358 ₄	0.316 ₂	0.264 ₇	0.300 ₈	0.293 ₆	0.341 ₂	0.304 ₃
OUR MODEL	0.330₃‡	0.351₃‡	0.428₅‡	0.452 ₁	0.360₄‡	0.362 ₅	0.360₅‡	0.370₂‡	0.370₁	0.378 ₂

Table 4: Results of present keyphrase prediction. Results shown in the upper part are directly cited from their corresponding papers. The subscript denotes the corresponding standard deviation (e.g., 0.330₃ indicates 0.330±0.003). ‡ indicates significant at $p < 0.01$ over SETTRANS with 1,000 bootstrap tests (Efron and Tibshirani, 1993).

revisit Figure 2, we re-assign “slot5” with “patch clustering”, “slot6” with “denoising” and no supervisory signal to “slot3”. Through the above process, we can convert the original target assignment m into a new one, where we use the conventional training objective (See Equation 2) adjusted with λ_{adp} (See Equation 3) to train our model.

6 Experiments

6.1 Setup

Datasets. We train various models and select the optimal parameters on the KP20k validation dataset (Meng et al., 2017). Then, we evaluate these models on five test datasets: Inspec (Hulth, 2003), NUS (Nguyen and Kan, 2007), Krapivin (Krapivin et al., 2009), SemEval (Kim et al., 2010), and KP20k. As implemented in (Yuan et al., 2020; Ye et al., 2021), we perform data preprocessing including tokenization, lowercasing, replacing all digits with the symbol $\langle digit \rangle$ and removing duplicated instances.

Baselines. We compare our WR-ONE2SET based model with the following baselines:

- **CATSEQ(R)** (Yuan et al., 2020). This is the most popular RNN-based model trained under the ONE2SEQ paradigm, formulating keyphrase generation as a sequence-to-sequence generation task.
- **CATSEQ** (Ye et al., 2021). It is also trained under the ONE2SEQ paradigm, but utilizing

Transformer as backbone.

- **UNIKEYPHRASE** (Wu et al., 2021). This is a large-scale pre-trained language model trained to extract and generate keyphrases jointly.
- **SETTRANS** (Ye et al., 2021). It is our most important baselines. Besides, we report the performance of three SETTRANS variants: **SETTRANS(w/o $\lambda_{pre}, \lambda_{abs}$)** that does not introduce any hyper-parameter to alleviate the negative effect of excessive \emptyset tokens, **SETTRANS(#SLOT=N)** that is equipped with $N/2$ and $N/2$ slots for present target keyphrases and absent ones, respectively, and **SETTRANS(w/ BATCHING)** which sorts all training instances in the increasing order of target keyphrase numbers and uses batch-wise randomized order to keep the padding length optimized.
- **PROMPTKP** (Wu et al., 2022). It firstly extracts keywords for automatic prompt construction, and then uses a mask-predict-based approach to generate the final absent keyphrase constrained by prompt.

Implementation Details. We use Transformer-base (Vaswani et al., 2017) to construct all models. During training, we choose the top 50,002 frequent tokens to form the predefined vocabulary. We use the Adam optimizer with a learning rate of 0.0001, and a batch size of 12. During inference, we employ greedy search to generate keyphrases.

Model	Inspec		NUS		Krapivin		SemEval		KP20k	
	$F_1@5$	$F_1@M$	$F_1@5$	$F_1@M$	$F_1@5$	$F_1@M$	$F_1@5$	$F_1@M$	$F_1@5$	$F_1@M$
<i>Existing Neural Keyphrase Generation Models</i>										
CATSEQ(R)	0.004	0.008	0.016	0.028	0.018	0.036	0.016	0.028	0.015	0.032
CATSEQ	0.010 ₂	0.019 ₄	0.028 ₂	0.048 ₂	0.032 ₁	0.060 ₄	0.020 ₅	0.023 ₃	0.023 ₁	0.046 ₁
UNIKEYPHRASE	0.012	0.022	0.026	0.037	—	—	0.022	0.029	0.032	0.058
SETTRANS	0.021 ₁	0.034 ₃	0.042 ₂	0.060 ₄	0.047 ₇	0.073 ₁₁	0.026 ₃	0.034 ₅	0.036 ₂	0.058 ₃
PROMPTKP	0.017	0.022	0.036	0.042	—	—	0.028	0.032	0.032	0.042
<i>Our Models</i>										
SETTRANS	0.020 ₃	0.031 ₄	0.044 ₅	0.061 ₈	0.050 ₂	0.073 ₁	0.030 ₂	0.037 ₁	0.038 ₁	0.059 ₁
SETTRANS(w/o $\lambda_{pre}, \lambda_{abs}$)	0.000 ₀	0.000 ₀	0.002 ₁	0.003 ₂	0.004 ₁	0.008 ₂	0.002 ₁	0.003 ₂	0.002 ₁	0.005 ₁
SETTRANS(#SLOT=12)	0.016 ₃	0.027 ₆	0.042 ₇	0.065 ₁₃	0.047 ₅	0.073 ₉	0.024 ₈	0.031 ₈	0.033 ₁	0.057 ₂
SETTRANS(#SLOT=16)	0.018 ₁	0.030 ₂	0.040 ₅	0.060 ₈	0.045 ₃	0.074 ₂	0.023 ₁	0.031 ₁	0.034 ₁	0.057 ₂
SETTRANS(#SLOT=24)	0.019 ₂	0.029 ₅	0.044 ₃	0.061 ₄	0.046 ₅	0.073 ₉	0.026 ₃	0.035 ₄	0.038 ₁	0.059 ₂
SETTRANS(#SLOT=28)	0.016 ₃	0.026 ₅	0.044 ₄	0.063 ₃	0.043 ₁	0.070 ₁	0.021 ₄	0.027 ₂	0.032 ₃	0.054 ₃
SETTRANS(w/ BATCHING)	0.023 ₁	0.030 ₄	0.050 ₅	0.067 ₆	0.049 ₃	0.059 ₉	0.034 ₄	0.038 ₆	0.045 ₂	0.058 ₂
OUR MODEL	0.025 ₂	0.034 ₄	0.057 ₅ ‡	0.071 ₃ ‡	0.057 ₁ ‡	0.074 ₂	0.040 ₃ ‡	0.043 ₅ ‡	0.050 ₁ ‡	0.064 ₂

Table 5: Results of absent keyphrase prediction.

To ensure a fair comparison with SETTRANS, we also set both slot numbers for present and absent keyphrases as 10, the target assignment step K as 2, λ_{pre} as 0.2 and λ_{abs} as 0.1, respectively. Particularly, we run all experiments three times with different random seeds and report the average results, so as to alleviate the impact of the instability of model training.

Evaluation Metrics. Following previous studies (Chen et al., 2020; Ye et al., 2021), we use macro averaged $F_1@5$ and $F_1@M$ to evaluate the quality of both present and absent keyphrases. When using $F_1@5$, if the prediction number is less than five, blank keyphrases are added to make the keyphrase number reach five. Particularly, we employ the Porter Stemmer⁵ to remove the identical stemmed keyphrases.

6.2 Main Results

Table 4 and Table 5 show the prediction results on present and absent keyphrases, respectively. We can draw the following conclusions:

First, our reproduced SETTRANS achieves comparable performance to Ye et al. (2021). Second, when removing both λ_{pre} and λ_{abs} from SETTRANS, its performance significantly drops, showing that the \emptyset token over-estimation severely limits its full potential. Third, we observe no improvements with different number of slots. Fourth, the commonly-used batching method for sequence

⁵<https://github.com/nltk/nltk/blob/develop/nltk/stem/porter.py>

Model	In-domain		Out-domain	
	$F_1@5$	$F_1@M$	$F_1@5$	$F_1@M$
<i>Present Keyphrase Prediction</i>				
OUR MODEL	0.370	0.378	0.370	0.384
w/o RE-ASSIGN	0.368	0.375	0.360	0.377
w/o WEIGHTING	0.365	0.393	0.340	0.374
RE-ASSIGN \Rightarrow RAND-ASSIGN	0.368	0.377	0.365	0.380
SETTRANS	0.359	0.392	0.336	0.373
<i>Absent Keyphrase Prediction</i>				
OUR MODEL	0.050	0.064	0.043	0.055
w/o RE-ASSIGN	0.047	0.062	0.042	0.053
w/o WEIGHTING	0.043	0.063	0.039	0.052
RE-ASSIGN \Rightarrow RAND-ASSIGN	0.048	0.063	0.042	0.053
SETTRANS	0.038	0.059	0.034	0.052

Table 6: Ablation study on keyphrase predictions.

generation is not beneficial for SETTRANS. Finally, our model significantly surpasses all baselines. These results strongly validate the effectiveness and generalization of our WR-ONE2SET paradigm.

6.3 Ablation Study

To better investigate the effectiveness of our proposed strategies on WR-ONE2SET, we report the performance of variants of our model on two test sets: 1) KP20k that is an **in-domain** one, and 2) the combination of Inspec, NUS, Krapivin and SemEval, which is **out-domain**. Here, we mainly consider three variants: 1) w/o RE-ASSIGN, which removes the target re-assignment mechanism from our model; and 2) w/o WEIGHTING. It discards the adaptive instance-level cost weighting strategy;

Model	In-domain	Out-domain
SETTRANS(w/o $\lambda_{pre}, \lambda_{abs}$)	0.747	0.809
SETTRANS	0.345	0.418
SETTRANS(w/ RE-ASSIGN)	0.301	0.386
SETTRANS(w/ WEIGHTING)	0.240	0.308
OUR MODEL	0.211	0.263

Table 7: The proportions of slots over-estimating \emptyset token.

and 3) RE-ASSIGN \Rightarrow RAND-ASSIGN. This variant randomly re-assigns targets to the slots in C_p .

As shown in Table 6, when removing the target re-assignment mechanism, we observe a performance degradation on keyphrase predictions. Likewise, the variant w/o WEIGHTING is obviously inferior to our model on most metrics. Therefore, we believe that our proposed strategies indeed benefit the generation of keyphrase set.

6.4 Analyses of \emptyset Token Over-Estimation

We also compare various models according to the proportion of slots over-estimating \emptyset tokens. Here, the proportion is the ratio between two slot numbers obtained from the whole training data: one is the number of slots that directly output \emptyset token via the vanilla prediction while generating correct keyphrases through the non- \emptyset prediction; and the other is the number of slots that generate correct keyphrases via the non- \emptyset prediction. Table 7 displays the results. The proportions of SETTRANS (w/o $\lambda_{pre}, \lambda_{abs}$) exceeds 70%, demonstrating the severe \emptyset token over-estimation of the ONE2SET paradigm. By comparison, the proportions of SETTRANS decrease, validating the effectiveness of fixed hyper-parameters λ_{pre} and λ_{abs} on alleviating the class imbalance of training data. Moreover, whether adaptive instance-level cost weighting strategy or target re-assignment mechanism is used alone, the proportions of SETTRANS can be further reduced. Particularly, our model achieves the lowest proportions, proving that our strategies can complement each other.

Besides, following Guo et al. (2017), we show the reliability diagram of SETTRANS and our model in Figure 3. It displays the relationship between the prediction confidence (the predicted probability of model) and the prediction accuracy within the confidence interval $[0, 0.2]$. Especially, the predictions within the confidence interval $[0, 0.2]$ account for 69.8% of all predictions. Please note that, if a model is well-calibrated, the gap




Figure 3: Reliability diagrams of SETTRANS and our model on the in-domain test set. “Gap” (areas marked with slash) denotes the difference between the prediction confidence and the prediction accuracy. Smaller gaps denote better calibrated outputs.

Model	In-domain			Out-domain		
	#Pre	#Abs	Dup	#Pre	#Abs	Dup
ORACLE	3.31	1.95	-	5.53	3.51	-
CATSEQ(R)	3.71	0.55	0.39	3.46	0.72	0.54
CATSEQ	4.64	1.16	0.26	4.34	1.28	0.38
SETTRANS	5.10	2.01	0.08	4.62	2.18	0.08
SETTRANS(w/ RE-ASSIGN)	5.40	2.64	0.10	4.83	2.72	0.09
SETTRANS(w/ WEIGHTING)	6.19	3.12	0.11	5.70	3.56	0.09
OURS MODEL	6.35	3.26	0.10	5.94	3.60	0.10

Table 8: Numbers and duplication ratios of predicted keyphrases on test datasets. “ORACLE” refers to the average number of target keyphrases.

between the confidence and the accuracy will be small. Overall, the gap of our model is less than that of SETTRANS, which demonstrates that our proposed strategies can calibrate the predictions with low confidence.

6.5 Diversity of Predicted Keyphrases

Follow previous studies (Chen et al., 2020; Ye et al., 2021), we report the average numbers of unique present and absent keyphrases, and the average duplication ratios of all predicted keyphrases, so as to investigate the ability of our model in generating diverse keyphrases,

Table 8 reports the results. As expected, our model generates more keyphrases than previous models and achieves a slightly higher duplication ratio than SETTRANS, however, significantly lower than ONE2SEQ-based models. Note that compared to SETTRANS, the $F_1@5$ and $F_1@M$ scores of our model are significantly improved, which demonstrates that our model performs much better on keyphrase generation.

Slot Type	SETTRANS	OUR MODEL
<i>Present Keyphrase Prediction</i>		
Slot(\emptyset)	61.2%	63.5% (+2.3%)
Slot(Target KP)	17.6%	21.8% (+4.2%)
Slot(\emptyset +Target KP)	21.2%	14.7% (-6.5%)
<i>Absent Keyphrase Prediction</i>		
Slot(\emptyset)	66.4%	69.6% (+3.2%)
Slot(Target KP)	9.3%	12.6% (+3.3%)
Slot(\emptyset +Target KP)	24.4%	17.8% (-6.6%)

Table 9: The proportions of slots with different target assignments for keyphrase predictions.

6.6 Analyses of Target Re-Assignment

Here, we still focus on the assigned targets during the model training mentioned at the beginning of Section 4 and conduct two types of analyses to better understand the effects of our mechanism.

First, we count the proportions of \emptyset tokens in assigned targets. Specially, the assigned \emptyset tokens accounts for 72.4% and 80.4% on present and absent keyphrase slots, respectively, but decrease to 67.6% and 72.3% in our model. Second, as implemented in Section 4, we still classify the instance slots into three categories and report their proportions in Table 9. We can find the proportions of Slots(\emptyset +KP), where slots are assigned with target keyphrase and \emptyset token at different iterations of model training, sharply decline. Besides, for each slot, we use the entropy of target assignment distribution to measure the stability of its supervisory signals. Furthermore, we average the entropy values of all slots to quantify the stability of supervisory signals for each instance. Consequently, we find that the entropy decreases in 68.2% of instances, increases in 26.3% of instances, and remain unchanged in 5.5% of instances. These results indicate that our target re-assignment mechanism indeed not only reduces excessive target \emptyset tokens, but also alleviates the instability of supervisory signals.

7 Conclusion

In this paper, we in-depth analyze the serious calibration errors of the ONE2SET paradigm and point out its underlying reasons. To deal with this issue, we then significantly extend the conventional ONE2SET into the WR-ONE2SET paradigm with an adaptive instance-level cost weighting strategy and a target re-assignment mechanism. Extensive experiments verify the effectiveness and generality of our extended paradigm.

In the future, we plan to further refine our WR-ONE2SET paradigm by considering the semantic relation between keyphrases. Besides, we will improve our model by introducing variational neural networks, which have been successfully applied in many NLP tasks (Zhang et al., 2016a,b; Su et al., 2018a,b; Liang et al., 2022). Finally, we will leverage the abundant knowledge from pre-trained models to further enhance our model.

Limitations

As mentioned above, serious \emptyset token over-estimation problem exists in ONE2SET paradigm, leading to a miscalibrated model. To solve this problem, we propose several strategies based on conventional ONE2SET using the same fixed hyper-parameters as Ye et al. (2021). However, hyper-parameter selection is a labor-intensive, manual, time-consuming process and affects generation performance deeply. Thus, our future work will focus on exploring a parameter-free method. Besides, despite achieving impressive performance, our WR-ONE2SET paradigm is only conducted based on the Transformer, so that it is essential to leverage the abundant knowledge from pre-trained models for better document modeling and keyphrase generation.

Acknowledgements

The project was supported by National Natural Science Foundation of China (No. 62276219, No. 62036004), Natural Science Foundation of Fujian Province of China (No. 2020J06001), and Youth Innovation Fund of Xiamen (No. 3502Z20206059). We also thank the reviewers for their insightful comments.

References

- Hou Pong Chan, Wang Chen, Lu Wang, and Irwin King. 2019. [Neural keyphrase generation via reinforcement learning with adaptive rewards](#). In *Proc. of ACL*, pages 2163–2174.
- Wang Chen, Hou Pong Chan, Piji Li, Lidong Bing, and Irwin King. 2019a. [An integrated approach for keyphrase generation via exploring the power of retrieval and extraction](#). In *Proc. of NAACL-HLT*, pages 2846–2856.
- Wang Chen, Hou Pong Chan, Piji Li, and Irwin King. 2020. [Exclusive hierarchical decoding for deep keyphrase generation](#). In *Proc. of ACL*, pages 1095–1105.

- Wang Chen, Yifan Gao, Jiani Zhang, Irwin King, and Michael R. Lyu. 2019b. [Title-guided encoding for keyphrase generation](#). In *Proc. of AAAI*, pages 6268–6275.
- Bradley Efron and Robert Tibshirani. 1993. *An Introduction to the Bootstrap*. Springer.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. [On calibration of modern neural networks](#). In *Proc. of ICML*, pages 1321–1330.
- Anette Hulth. 2003. [Improved automatic keyword extraction given more linguistic knowledge](#). In *Proc. of EMNLP*, pages 216–223.
- Anette Hulth and Beáta Megyesi. 2006. [A study on automatically extracted keywords in text categorization](#). In *Proc. of ACL*, pages 537–544.
- Su Nam Kim, Olena Medelyan, Min-Yen Kan, and Timothy Baldwin. 2010. [Semeval-2010 task 5 : Automatic keyphrase extraction from scientific articles](#). In *Proc. of SemEval@ACL*, pages 21–26.
- Mikalai Krapivin, Aliaksandr Autaeu, and Maurizio Marchese. 2009. Large dataset for keyphrases extraction.
- Harold W. Kuhn. 2010. [The hungarian method for the assignment problem](#). In *50 Years of Integer Programming 1958-2008 - From the Early Years to the State-of-the-Art*, pages 29–47.
- Aviral Kumar and Sunita Sarawagi. 2019. [Calibration of encoder decoder models for neural machine translation](#). In *Proc. of ICLR Debugging Machine Learning Models Workshop*.
- Yunlong Liang, Fandong Meng, Chulun Zhou, Jinan Xu, Yufeng Chen, Jinsong Su, and Jie Zhou. 2022. [A variational hierarchical model for neural cross-lingual summarization](#). In *Proc. of ACL*, pages 2088–2099.
- Rui Meng, Sanqiang Zhao, Shuguang Han, Daqing He, Peter Brusilovsky, and Yu Chi. 2017. [Deep keyphrase generation](#). In *Proc. of ACL*, pages 582–592.
- Rada Mihalcea and Paul Tarau. 2004. [TextRank: Bringing order into text](#). In *Proc. of EMNLP*, pages 404–411.
- Thuy Dung Nguyen and Min-Yen Kan. 2007. [Keyphrase extraction in scientific publications](#). In *Proc. of ICADL*, pages 317–326.
- Jinsong Su, Shan Wu, Deyi Xiong, Yaojie Lu, Xianpei Han, and Biao Zhang. 2018a. [Variational recurrent neural machine translation](#). In *Proc. of AAAI*, pages 5488–5495.
- Jinsong Su, Shan Wu, Biao Zhang, Changxing Wu, Yue Qin, and Deyi Xiong. 2018b. [A neural generative autoencoder for bilingual word embeddings](#). *Inf. Sci.*, 424:287–300.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Proc. of NIPS*, pages 5998–6008.
- Xiaojun Wan and Jianguo Xiao. 2008. [Single document keyphrase extraction using neighborhood knowledge](#). In *Proc. of AAAI*, pages 855–860.
- Lu Wang and Claire Cardie. 2013. [Domain-independent abstract generation for focused meeting summarization](#). In *Proc. of ACL*, pages 1395–1405.
- Huanqin Wu, Wei Liu, Lei Li, Dan Nie, Tao Chen, Feng Zhang, and Di Wang. 2021. [Unikeyphrase: A unified extraction and generation framework for keyphrase prediction](#). In *Proc. of ACL Findings*, pages 825–835.
- Huanqin Wu, Baijiaxin Ma, Wei Liu, Tao Chen, and Dan Nie. 2022. [Fast and constrained absent keyphrase generation by prompt-based learning](#). In *Proc. of AAAI*, pages 11495–11503.
- Jiacheng Ye, Tao Gui, Yichao Luo, Yige Xu, and Qi Zhang. 2021. [One2Set: Generating diverse keyphrases as a set](#). In *Proc. of ACL*, pages 4598–4608.
- Xingdi Yuan, Tong Wang, Rui Meng, Khushboo Thaker, Peter Brusilovsky, Daqing He, and Adam Trischler. 2020. [One size does not fit all: Generating and evaluating variable number of keyphrases](#). In *Proc. of ACL*, pages 7961–7975.
- Biao Zhang, Deyi Xiong, Jinsong Su, Hong Duan, and Min Zhang. 2016a. [Variational neural machine translation](#). In *Proc. of EMNLP*, pages 521–530.
- Biao Zhang, Deyi Xiong, Jinsong Su, Qun Liu, Rongrong Ji, Hong Duan, and Min Zhang. 2016b. [Variational neural discourse relation recognizer](#). In *Proc. of EMNLP*, pages 382–391.

A Appendix

A.1 Example

Input Document: an image topic model for image denoising. topic model is a powerful tool for the basic document or image processing tasks. in this study, we introduce a novel image topic model, called latent patch model (lpm), which is a generative bayesian model and assumes that the image and pixels are connected by a latent patch layer. based on the lpm, we further propose an image denoising algorithm namely multiple estimate lpm (melpm). unlike other works, the proposed denoising framework is totally implemented on the latent patch layer, and it is effective for both gaussian white noises and impulse noises. experimental results demonstrate that lpm performs well in representing images...

Keyphrases: topic model; denoising; patch clustering; semantic learning

Table 10: An example of keyphrase generation. The underlined phrases are *present keyphrases* that appear in the document, and other phrases are *absent keyphrases* that do not match any contiguous subsequence of the document.

A.2 Formal Definitions

Symbol	Definition
\emptyset	A special token representing no corresponding keyphrase.
N	The predefined number of slots generating keyphrases or \emptyset tokens in parallel.
K	The predefined number of tokens generated from each slot for the conventional target assignment mechanism.
\mathbf{P}^i	The predicted probability distributions of the slot i .
$\hat{p}^i(*)$	The predicted probability of a keyphrase from the slot i using teacher forcing. Specially, the j -th token predictive probability of $\hat{p}^i(*)$ is denoted as $\hat{p}_j^i(*)$.
$M(N)$	The set of all N -length target index permutations.
m	The optimal permutation of $M(N)$. It can be considered as a mapping function from the slot i to the target index $m(i)$. Particularly, we use m^p and m^a to denote the optimal permutations for present and absent keyphrases, respectively.
$y^{m(i)}$	The assigned target of the slot i .
$\lambda_{pre}, \lambda_{abs}$	Two predefined hyper-parameters used to reduce the negative effect of excessive \emptyset tokens for present and absent keyphrase predictions, respectively.
λ_{adp}	The degree of \emptyset token over-estimation for each instance, which is leveraged to dynamically scale the losses corresponding to \emptyset tokens in our paradigm.
$\mathcal{L}(\theta)$	The training loss of the whole model with parameters θ . Moreover, we use $\mathcal{L}^p(\theta, z)$ to denote the present keyphrase training loss on the assigned target z . The absent keyphrase training loss $\mathcal{L}^a(\theta, z)$ is defined in a similar way.
$\hat{y}_{:K}^i$	The first K tokens of the prediction from the slot i via the vanilla prediction.
$\bar{y}_{:K}^i$	The first K tokens of the prediction from the slot i through the non- \emptyset prediction, where \emptyset token is removed from the prediction vocabulary.
C_{\emptyset}	The set of slots, each of which is assigned with a keyphrase as supervisory signal.
C_p	The set of potential slots, where each slot has the potential to generate a fresh keyphrase, boosting the performance of the model.
C_u	The unimportant slot set, where each slot has no effect on the model performance.

Table 11: Formal Definitions