

# From Lazy to Prolific: Tackling Missing Labels in Open Vocabulary Extreme Classification by Positive-Unlabeled Sequence Learning

Ranran Haoran Zhang<sup>1</sup>, Bensu Uçar<sup>2</sup>, Soumik Dey<sup>2</sup>, Hansi Wu<sup>2</sup>, Binbin Li<sup>2</sup>, Rui Zhang<sup>1</sup>

<sup>1</sup> Penn State University

<sup>2</sup> eBay Inc

## Abstract

Open-vocabulary Extreme Multi-label Classification (OXMC) extends traditional XMC by allowing prediction beyond an extremely large, predefined label set (typically  $10^3$  to  $10^{12}$  labels), addressing the dynamic nature of real-world labeling tasks. However, self-selection bias in data annotation leads to significant missing labels in both training and test data, particularly for less popular inputs. This creates two critical challenges: generation models learn to be “lazy” by under-generating labels, and evaluation becomes unreliable due to insufficient annotation in the test set. In this work, we introduce Positive-Unlabeled Sequence Learning (PUSL), which reframes OXMC as an infinite keyphrase generation task, addressing the generation model’s laziness. Additionally, we propose to adopt a suite of evaluation metrics,  $F1@O$  and newly proposed  $B@k$ , to reliably assess OXMC models with incomplete ground truths. In a highly imbalanced e-commerce dataset with substantial missing labels, PUSL generates 30% more unique labels, and 72% of its predictions align with actual user queries. On the less skewed EURLex-4.3k dataset, PUSL demonstrates superior F1 scores, especially as label counts increase from 15 to 30. Our approach effectively tackles both the modeling and evaluation challenges in OXMC with missing labels. Our data split was released in huggingface datasets<sup>1</sup>

## 1 Introduction

Extreme multi-label classification (XMC) is a challenging and critical task in natural language processing, with wide-ranging applications such as e-commerce query recommendations (Jain et al., 2019; Zhou et al., 2019; Song et al., 2021; Zhang et al., 2023; Mishra et al., 2024), tag recommendation for Wikipedia articles (Dekel and Shamir,

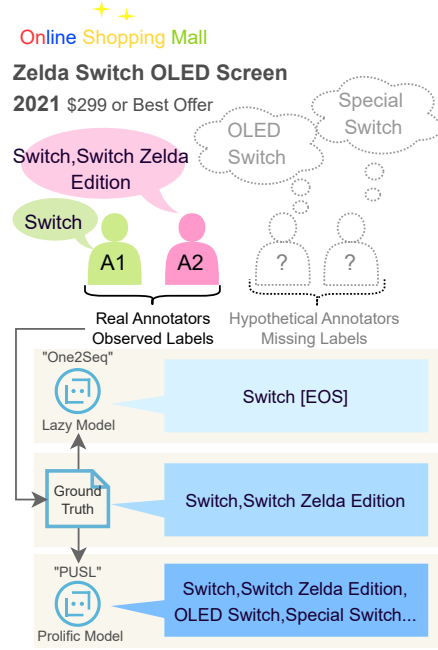


Figure 1: Real annotators often provide limited labels (e.g., Switch, Switch Zelda Edition), while many potential labels remain uncaptured (e.g., OLED Switch, Special Switch). This gap between observed and expected labels misleads generation models to be lazy by prematurely terminating label generation. Our proposed PUSL resolves the model laziness problem and can learn from incomplete ground truth.

2010), legal document (Chalkidis et al., 2019), stackoverflow questions (Kang et al., 2021) and social media posts (Wang et al., 2019; Yu et al., 2023). XMC aims to assign appropriate labels to each input from a vast pool of possibilities, often numbering from thousands to trillions.

Open-vocabulary Extreme Multi-label Classification (OXMC) bridges the traditional XMC and well-studied keyphrase generation by introducing the generation models (Bhatia et al., 2015; Zhang et al., 2016; Prabhu et al., 2018; Wang et al., 2019; Chang et al., 2021) into the XMC framework (Simig et al., 2022). While XMC assigns labels

<sup>1</sup><https://huggingface.co/datasets/windchimeran/pusl>

from a vast, predefined set, OXMC recognizes that label sets are continuously evolving, with new labels emerging over time. This shift acknowledges that predefined label sets, no matter how extensive, may not capture all possible ways to categorize input instances, especially in rapidly changing domains such as e-commerce and social media. Due to the larger scale of the datasets, successfully addressing the challenges posed by OXMC could also benefit the general keyphrase generation area.

However, our research reveals that the prevalence of missing labels in OXMC training data leads to “lazy” generation models which predict only a few labels per input. This problem stems from the self-selection bias in the data curation processes of OXMC, as illustrated in Figure 1. Annotators tend to selectively tag data points, with popular inputs receiving more annotations. This self-selection bias causes insufficient annotators and labels on most data. In our dataset (Ads-XMC), 89% of data points have fewer than five labels, with only 15% of these ever having sufficient annotators. In training, these cause the models to prematurely stop label generation before producing a sufficient number of labels. Moreover, existing evaluation metrics fail to accurately reflect this problem in test data, further exacerbating the issue.

To address the generation model’s “laziness” problem, we introduce Positive-Unlabeled Sequence Learning (PUSL) to foster prolific and accurate label prediction. PUSL treats observed labels as positive, and treats potential but unobserved labels as unlabeled, focusing on positive data to infer a hypothetically infinite sequence of labels. This method views existing labels as part of a larger conceivable set, aiming to mirror exhaustive annotations and mitigate bias from training data. By expanding beyond initially observed labels, PUSL overcomes gaps from insufficient annotations, increasing the diversity of labels generated across all inputs. To further enhance PUSL’s capability in generating diverse labels, we implement a post-training phase that leverages both original diverse training data and augmented data created through rejection sampling, exposing the model to a more comprehensive label set.

To address the challenges in evaluating OXMC models with incomplete ground truth due to missing annotations, we propose two metrics:  $F1@O$  and  $B@k$ . Traditional metrics like  $F1@k$  tend to favor lazy models when faced with incomplete test data. Our proposed  $F1@O$  mitigates this bias by

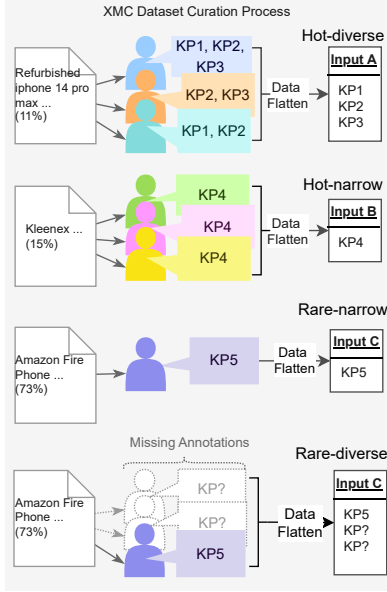
adapting to the variable number of available ground truth labels, avoiding penalties for prolific models. Additionally,  $B@k$  addresses the limitations of existing metrics by penalizing under-generation while fairly assessing predictions beyond available ground truth. Together, these metrics provide a more faithful assessment of OXMC models in real-world scenarios with incomplete annotations, effectively distinguishing between lazy and prolific prediction behaviors.

Our experimental results demonstrate PUSL’s effectiveness and applicability across various domains. In the highly-imbalanced Ads-XMC dataset with substantial missing labels, PUSL generates 30% more unique labels per item, and 72% of its predicted keyphrases align with actual user queries. In the EURLex4.3k dataset, with an average of 15 keyphrases per data point, PUSL shows promising results over all baselines in  $F1@O$  and  $B@k$ , particularly as  $k$  increases from 15 to 30. Our contributions include:


- **Data Imbalance Analysis:** We identify self-selection bias in OXMC datasets as the cause of substantial missing labels.
- **Model Solution:** We recast OXMC as Positive Unlabeled Sequence Learning, and our experiment demonstrates its effectiveness to foster prolific and accurate label prediction.
- **Evaluation Solution:** We propose  $F1@O$  and  $B@k$  metrics to reliably evaluate models despite incomplete ground-truth labels.

## 2 Related Work


**Missing Labels in XMC.** The large label set in XMC challenges exhaustive annotation, resulting in two types of missing labels: insufficient labels in the training set (Schultheis and Babbar, 2021; Schultheis et al., 2022, 2023) and unseen labels in the test set (Gupta et al., 2021). Existing approaches like GROOV (Simig et al., 2022) address unseen labels but tend to predict fewer, while XL-Gen (Jung et al., 2023) is limited to seen labels due to its pre-clustering process. Previous studies attribute missing labels to annotator subjectivity and heedlessness (Sterckx et al., 2016; Lei et al., 2021). Wei and Li (2019) theoretically analyze the impact of missing labels under the Missing Completely at Random (MCAR) assumption (Rubin, 1976). By contrast, our investigation reveals that the missing label mechanism in XMC deviates



(a) The XMC dataset curation process. KP stand for keyphrase. Popular inputs like iPhone or Kleenex receive more annotations, while others are neglected.



(b) Distribution of numbers of items by unique keyphrases in Ads-XMC.



(c) Frequency distribution of the number of keyphrases per input for various XMC datasets. Higher CVs indicating greater disparity in the number of keyphrases per input.

Figure 2: Analysis of Missing Labels in OXMC.

from the commonly assumed MCAR, because it stems from the unequal distribution of annotations from self-selection bias (Marlin et al., 2012). We introduce an improved generation model capable of addressing both unseen and insufficient labels by explicitly modeling this distribution, empowering XMC models for the complexities of real-world data.

**Positive Unlabeled Learning.** Traditional PU learning trains a binary classifier to differentiate between positive and negative examples without labeled negatives in the dataset (Lee and Liu, 2003; Liu et al., 2003; Elkan and Noto, 2008; Kiryo et al., 2017). This has evolved into Multi-Positive and Unlabeled learning for multi-class scenarios, where training data includes labeled examples from several positive classes and unlabeled examples which may be positive or a single negative class (Xu et al., 2017). In XMC, the large label space precludes simple conversion to multiple binary classifications, while our proposed PUSL leverages open vocabulary generation models to make this task feasible. Moreover, our investigation reveals that the missing label mechanism in XMC deviates from the commonly assumed MCAR condition in PU learning, aligning more with Selected At Random PU Learning (Bekker et al., 2019; Bekker and Davis, 2020), indicating systematic patterns or biases in

label absence.

### 3 Analysis on Missing Label in OXMC

To investigate the effect of self-selection bias on the label distribution, we analyzed an e-commerce dataset, Ads-XMC. Unlike most public OXMC datasets, Ads-XMC retains crucial information on user interaction frequencies (i.e., annotation frequency for each input), making it uniquely suitable for this analysis. The dataset consists of item-keyphrase pairs derived from user interactions, along with interaction frequencies. During the OXMC data curation process, as shown in Figure 2a, the raw data is transformed into a flattened dataset by deduplicating labels and retaining only a single instance of each unique label per item, aligning with the format used in previous public datasets. Thus, the number of keyphrases is bounded by the total user interaction, if no duplication.

We classified data based on item popularity (hot or rare) and label diversity (diverse or narrow) using a threshold of five interactions or labels, identified as a significant breakpoint for user interest. Hot items have  $\geq 5$  interactions (top 26%), while diverse items have  $\geq 5$  unique labels. This categorization data samples in four groups: (1) **Hot-diverse**: High interaction with a wide variety of labels; (2) **Hot-narrow**: Many interactions but a

limited range of labels; (3) **Rare-narrow**: Few interactions and a small set of labels; and (4) **Rare-diverse**: Potentially broad label spectrum if more interactions occurred.

As shown in Figure 2b, Rare-narrow items cover 73.7% of the dataset, while Hot-narrow items, despite having few unique keyphrases per item, cover only 15.2%. Annotators’ personal choices lead to an uneven distribution of attention, similar to the 80/20 rule: a small number of widely recognized items attract most focus, while the majority remain less popular and under-annotated. This results in missing labels, affecting both training and test data in OXMC datasets, and posing significant challenges for model training and evaluation.

The issue of *self-selection bias* (unequal distribution of annotations) is pervasive in OXMC datasets, though most lack the frequency data to quantify this imbalance. For example, the *Barack Obama* Wikipedia page likely receives more revisions and annotations than the less-known *Andrew Dunlop* page, resulting in a more comprehensive tag set for the former. Similarly, on music streaming platforms, tracks by well-known artists like *Beyoncé* are tagged with more genres and moods than those by lesser-known musicians.

Although there is no annotation frequency information in other datasets, we quantify the label missingness by the coefficient of variation (CV) (Brown, 1998):  $CV = \mu/\sigma$ , where  $\mu$  is the mean number of keyphrases and  $\sigma$  is the standard deviation. A higher CV indicates more imbalance. Figure 2c shows that OXMC datasets have different levels of keyphrase imbalance, with Ads-XMC being the most imbalanced. The Ads-XMC hot split has a distribution similar to AmazonCat and a lower CV than the Ads-XMC all split, demonstrating different degrees of missing labels. Key takeaways from the analysis:

- Missing labels result from self-selection bias, with few items receiving most labels.
- Both training and test data suffer from missing labels, making model training and evaluation challenging.
- The imbalance and severity of missing labels varies across datasets, measurable by the coefficient of variation of keyphrases per input.

## 4 Method

We present our PUSL approach to address the challenges of missing labels in OXMC. Additionally,

we propose evaluation metrics that account for missing labels in the test data, enabling a more accurate assessment of model performance.

### 4.1 Task: OXMC as Positive-Unlabeled Sequence Generation

Given a textual input  $x$ , XMC aims to predict a set of labels  $\mathcal{Y} = \{y_1, y_2, \dots, y_{|\mathcal{Y}|}\}$  from an extremely large label space. Each label  $y_j = [w_1, w_2, \dots, w_{|y_j|}]$  is a sequence of words (i.e., keyphrase). We formulate OXMC as a Positive-Unlabeled Sequence Generation problem:

$$\mathcal{D} = \{(x_i, \mathcal{Y}_i^p)\}_{i=1}^n \cup \{(x_i, \mathcal{Y}_i^u)\}_{i=1}^n$$

$$\mathcal{D}_{train} \subset \{(x_i, \mathcal{Y}_i^p)\}_{i=1}^n, \mathcal{D}_{test} \subset \{(x_i, \mathcal{Y}_i^p)\}_{i=1}^n \quad (1)$$

where  $\mathcal{Y}_i^p$  is the observed label set and  $\mathcal{Y}_i^u$  is the unobserved label set for document  $x_i$ .

Real-world XMC datasets typically only include observed keyphrases due to annotation challenges. The goal is to learn a top- $k$  keyphrase generator  $f : (x, k) \rightarrow \mathcal{Y}$  that maps  $x$  to  $\{y_1, y_2, \dots, y_k\}$ , where  $k$  is a predefined requirement (e.g., determined by the constraints of the user interface or website), but may exceed the number of observed keyphrases in most labeled data.

### 4.2 Background: Keyphrase Generation Biases

Existing keyphrase generation paradigms face biases when applied to OXMC tasks due to incomplete training data. We discuss two main paradigms and their associated biases as shown in Figure 3:

**One2Seq.** This paradigm treats keyphrase generation as a single sequence generation task (Chan et al., 2019; Yuan et al., 2020). The model learns to generate all keyphrases as a concatenated sequence, separated by delimiters, and terminated by an end-of-sequence (EOS) token. **GROOV**, the state-of-the-art generation One2Seq model for OXMC, employs a label trie and multi-softmax loss for unordered training (Simig et al., 2022). **Early Termination Bias.** When trained on incomplete XMC data, One2Seq models tend to prematurely generate the EOS token, resulting in shorter keyphrase sequences. This bias stems from the model learning to terminate generation based on the limited number of observed keyphrases in the training data, rather than the true, exhaustive set of relevant keyphrases.



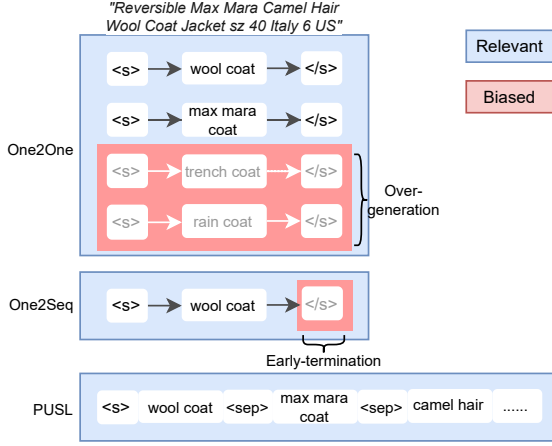


Figure 3: Comparison of biases in keyphrase generation models: One2Seq’s early termination, One2One’s over-generation, and PUSL.

**One2One.** In this paradigm, each training instance consists of an input document paired with only one keyphrase (Meng et al., 2017; Chen et al., 2019a,b). During inference, the model generates top- $k$  candidate keyphrases through beam search.

**Over-generation Bias.** One2One models often exhibit over-generation bias, especially when producing more keyphrases than the input document can support. This bias arises from the beam search decoding process (Kang et al., 2021), which can generate keyphrases without considering their absolute probabilities, potentially including low-probability, irrelevant candidates. When the requested number of keyphrases exceeds the relevant ones, the model may produce low-quality or irrelevant keyphrases to meet the quota, leading to their inclusion in the final output.

### 4.3 Model Solution: Positive-Unlabeled Sequence Learning

**Training.** To tackle the missing label problem, our PUSL model treats the training target keyphrases as an infinite sequence. Unlike standard sequence-to-sequence models that use an end-of-sequence (EOS) token, PUSL omits this token, allowing the model to generate keyphrases continuously without a defined endpoint.

**Post Training.** PUSL exhibits a repetition problem when generating keyphrases beyond the input’s natural limit, similar to cognitive overload. This issue occurs when the model is pushed to produce more keyphrases than the input supports. To address its repetition problem and enhance the keyphrase generation, we implement a post-training phase using two data sources: (1) Original

### Algorithm 1 Top- $k$ PUSL Decoding

---

**Require:**  $\text{token}_{1:n-1}$ , logits,  $k$

- 1:  $\text{eok\_count} \leftarrow \text{count}(\text{tokens}_{1:n-1}, \text{EOK})$
- 2:  $\text{completed\_sequences} \leftarrow \text{eos\_counts} \geq k$
- 3:  $\text{completed\_keyphrase} \leftarrow \text{token}_{n-1} == \text{EOK}$
- 4:  $\triangleright$  Force the model to decode a new keyphrase after generating EOK  $\triangleleft$
- 5:  $\text{logits}[\text{completed\_keyphrase}, :] \leftarrow -\infty$
- 6:  $\text{logits}[\text{completed\_keyphrase}, \text{EOK}] \leftarrow 0$
- 7:  $\triangleright$  Terminate generation after generating  $k$  EOK  $\triangleleft$
- 8:  $\text{logits}[\text{completed\_sequences}, :] \leftarrow -\infty$
- 9:  $\text{logits}[\text{completed\_sequences}, \text{EOS}] \leftarrow 0$
- 10: **return** logits

---

**Train-Diverse data:** We select instances from the original training set with an above-average number of keyphrases, focusing on richer annotations. (2) **Augmented data from rejection sampling:** We use PUSL to generate additional keyphrases for instances with few keyphrases, then apply rejection sampling to select those where the total number (original + generated) exceeds the dataset mean after deduplication. This creates a synthetic dataset with a higher density of plausible keyphrases.

**Inference.** PUSL employs an auto-regressive decoding strategy for keyphrase generation. Although designed for infinite generation, we extract only the top- $k$  results by adjusting the logits processor during decoding. Algorithm 1 outlines this process, which uses Begin-of-Keyphrase (BOK) and End-of-Keyphrase (EOK) tokens to delineate boundaries. This approach, as opposed to using a single separator (SEP), simplifies the decoding logit processor implementation and reduces the occurrence of empty keyphrases.

### 4.4 Evaluation Solution: Faithful Metrics for Incomplete Ground Truth

Evaluating OXMC models is challenging due to missing labels in the test data, which contains data points with fewer keyphrases than expected. We begin by discussing the limitations of existing metrics in the presence of missing labels and then propose new metrics for more accurate comparison.

Traditional metrics like Precision@ $k$  (P@ $k$ ) use a fixed  $k$  across all test data points, which can unfairly favor “lazy” models that generate fewer keyphrases, especially in datasets with many missing labels. As shown in Figure 4, even if both lazy and prolific models are perfect predictors (only generating keyphrases in the ground truth), the lazy model achieves higher precision scores due to smaller denominators, while the prolific model is

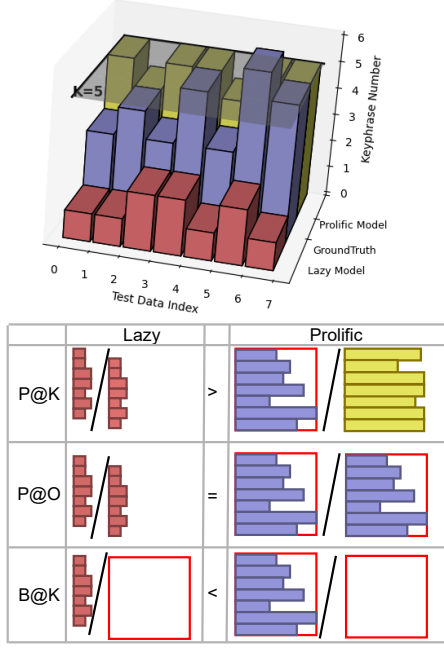


Figure 4: Comparison of evaluation metrics for lazy and prolific models under incomplete ground truth. Ground truth beyond  $K$  are truncated in the denominator. Traditional  $P@K$  favors lazy models even when all predictions are relevant. Proposed  $P@O$  equalizes performance between lazy and prolific models, while  $B@K$  penalizes under-generation. These metrics provide a more faithful evaluation with missing ground truth.

penalized for generating more keyphrases than are present in the ground truth.

To address this issue, we adopt  $F1@O$  (Yuan et al., 2020) which is the F1 score of the following:<sup>2</sup>

$$P@O = \frac{|\hat{\mathcal{Y}}_{:O} \cap \mathcal{Y}|}{|\hat{\mathcal{Y}}_{:O}|} \quad R@O = \frac{|\hat{\mathcal{Y}}_{:O} \cap \mathcal{Y}|}{|\mathcal{Y}|} \quad (2)$$

where  $|\cdot|$  denotes the cardinality of a set,  $\mathcal{Y}$  is the ground truth,  $O$  is the number ground truth labels,  $\hat{\mathcal{Y}}_{:O}$  is model’s top- $O$  predicted keyphrases. As shown in Figure 4,  $F1@O$  adjusts to the varying number of ground truth labels per data point, equalizing the performance of perfect lazy and prolific models.

Furthermore, in practice,  $F1@O$  may still favor imperfect lazy models, as predicting more relevant keyphrases is inherently more challenging, and prolific models may be unfairly penalized by missing ground truth. To provide a more faithful evaluation that considers both fairness and the user-defined  $k$ , we propose a new metric BudgetAccuracy@ $k$

( $B@k$ ):

$$B@k = \frac{|\hat{\mathcal{Y}}_{:k} \cap \mathcal{Y}|}{k} \quad (3)$$

where  $\hat{\mathcal{Y}}_{:k}$  represents the model’s top- $k$  predicted keyphrases and  $\mathcal{Y}$  is the ground truth.  $B@k$  addresses variable ground truth sizes and under-generation by models: (1) It penalizes models for generating fewer than  $k$  keyphrases by treating missing predictions as empty. (2) It avoids penalizing models for generating more keyphrases than available in the ground truth.  $B@k$  measures the model’s ability to generate the desired number of keyphrases while maintaining relevance to the ground truth. The upper bound of  $B@k$  is determined by  $\min(|\mathcal{Y}|, k)/k$ , reflecting the limitations of the available ground truth. This makes  $B@k$  particularly suitable for evaluating OXMC tasks where the ground truth may be incomplete.

This suite of metrics addresses challenges posed by incomplete ground truth and varying model behaviors in OXMC tasks. By using these metrics in combination, they enable a more faithful comparison between lazy and prolific models, accurately reflecting a model’s ability to generate comprehensive and relevant keyphrase sets.

## 5 Experimental Settings

### 5.1 Datasets and Diverse / Narrow Splits

We evaluate our approach on two public datasets commonly used in extreme multi-label classification research: **Eurlex 4.3k** (Chalkidis et al., 2019) and **AmazonCat** (McAuley and Leskovec, 2013). We also include **Ads-XMC** derived from search logs, where buyers act as annotators by clicking on items in the search results page for a given query.

Traditional XMC datasets were split to ensure each label appeared at least once in training (Gupta et al., 2021), rather than maintaining independent and identically distributed (IID) samples. As OXMC models now handle unseen labels (Simig et al., 2022), we introduce a new splitting procedure that prioritizes IID evaluation while enabling analysis across different keyphrase densities. Our procedure (Algorithm 2) re-splits the data into **Test-Narrow** ( $|\mathcal{Y}| \leq 2\mu$ ) and **Test-Diverse** ( $|\mathcal{Y}| > 2\mu$ ) based on the mean number of keyphrases  $\mu$ , allowing systematic evaluation of model performance on inputs with both few and many associated keyphrases. Dataset statistics and detailed splitting process are provided in Appendix A.

<sup>2</sup>Precision@ $O$  is identical to R-Precision@ $k$  (Schütze et al., 2008; Chalkidis et al., 2021)

Dataset	Model	Test-Narrow ( $ \mathcal{Y}  \leq 2\mu$ )					Test-Diverse ( $ \mathcal{Y}  > 2\mu$ )				
		#K@ $\mu$	#K@ $2\mu$	B@ $\mu$	B@ $2\mu$	F1@ $\mathcal{O}$	#K@ $\mu$	#K@ $2\mu$	B@ $\mu$	B@ $2\mu$	F1@ $\mathcal{O}$
EURLex4.3k ( $\mu = 15$ )	One2One	$\mu$	$2\mu$	35.15	22.73	36.68	$\mu$	$2\mu$	25.98	16.63	14.23
	GROOV	7.31	7.56	24.09	12.14	36.94	13.57	17.33	69.06	41.42	43.06
	One2Seq	14.05	<b>22.36</b>	52.01	33.32	57.77	14.88	<b>28.98</b>	77.05	80.02	74.40
	PUSL	<b>14.56</b>	22.02	54.75	33.64	58.81	<b>14.95</b>	28.78	88.56	82.29	80.28
AmazonCat ( $\mu = 5$ )	One2One	$\mu$	$2\mu$	48.37	38.58	48.06	$\mu$	$2\mu$	70.90	60.26	52.64
	GROOV	2.97	2.99	48.33	24.34	62.54	3.58	3.62	63.79	32.26	35.92
	One2Seq	4.71	8.83	38.74	30.63	37.50	4.97	9.90	64.28	57.57	55.58
	PUSL	<b>4.91</b>	<b>9.49</b>	47.67	38.71	47.12	<b>5.00</b>	<b>9.97</b>	74.82	69.16	61.64
Ads-XMC ( $\mu = 3$ )	One2One	$\mu$	$2\mu$	34.78	24.47	35.15	$\mu$	$2\mu$	59.32	48.27	34.37
	GROOV	1.03	1.03	18.72	9.36	28.95	1.07	1.07	26.01	13.00	12.56
	One2Seq	1.89	2.17	25.52	13.75	31.26	1.85	2.10	41.43	22.73	20.44
	PUSL	2.46	3.31	34.50	19.48	38.54	2.46	3.23	56.14	33.79	27.94
	PUSL+PT	<b>2.78</b>	<b>4.27</b>	34.73	21.03	36.47	<b>2.76</b>	<b>4.17</b>	62.52	41.21	30.09

Table 1: Evaluation results for various OXMC datasets. We present F1@ $\mathcal{O}$ , B@ $k$ , and number of unique keyphrases (#K@ $k$ ) where  $k$  is chosen as  $\mu$  or  $2\mu$ . Lighter color means higher results. Among all One2Seq family models, PUSL demonstrates the highest prolificacy across all datasets. One2One shows high performance in Ads-XMC, where the keyphrase set is typically narrow but performs poorly on the diverse EURLex4.3k dataset.

## 5.2 Baselines and Implementation Details

We compare PUSL with three state-of-the-art keyphrase generation baselines: (1) self-terminated generation models: **One2Seq** (Chan et al., 2019; Yuan et al., 2020) and **GROOV** (Simig et al., 2022), and (2) top- $k$  generation models: **One2One** (Meng et al., 2017; Chen et al., 2019a,b). We focus on generation models since traditional classification models cannot predict beyond a predefined label set (Simig et al., 2022). For fair comparisons, all models use the same t5-base architecture (Raffel et al., 2020) and identical hyperparameters. We only apply post-training on PUSL for Ads-XMC dataset, due to the high repetition rate. Hyperparameters and post-training ablation studies are provided in Appendix B and C.

## 6 Results and Analysis

We compare PUSL with state-of-the-art baselines across various datasets and metrics, focusing on incomplete ground truth (Section 6.1), early-termination bias in self-terminated models (Section 6.2), and over-generation bias in top- $k$  models (Section 6.3). Our analysis includes standard metrics, human evaluations, and an exhaustive annotation experiment (Section 6.6), providing comprehensive insights into model performance and dataset characteristics.

### 6.1 Evaluation using Incomplete Ground Truth

Table 1 presents our evaluation results across different datasets and models, highlighting the chal-

lenges of evaluation with incomplete ground truth.

**Evaluate with More Ground Truth: Test-Diverse vs Test-Narrow.** Our results show a significant performance gap between the Test-Diverse and Test-Narrow subsets across all models. Test-Diverse, with more ground truth labels, generally yields higher BudgetAccuracy and F1 scores, indicating sufficient annotation. In contrast, Test-Narrow likely suffers from missing ground truth, highlighting the impact of incomplete labels on evaluation metrics.

**Ask Models to be More Prolific: B@ $\mu$  vs B@ $2\mu$ .** Comparing B@ $\mu$  and B@ $2\mu$  evaluates models’ ability to generate relevant keyphrases beyond the training data average. In Test-Diverse, prolific models show minimal performance decline from B@ $\mu$  to B@ $2\mu$ , while Test-Narrow exhibits sharp drops due to limited ground truth. The performance gap between B@ $\mu$  and B@ $2\mu$  also reflects ground truth distribution: a significant drop in Ads-XMC Test-Narrow versus a smaller decline in AmazonCat Test-Narrow suggests a smoother distribution in the latter, as seen in Figure 2c.

**Toward More Faithful Evaluation: F1 vs BudgetAccuracy.** In Test-Narrow scenarios with sparse ground truth, F1 scores can misleadingly favor “lazy” models (Section 4.4). For instance, in AmazonCat-Test-Narrow, GROOV achieves the highest F1@ $\mathcal{O}$  (62.54) by generating fewer keyphrases (2.99 on average when asked for 10), coincidentally matching the sparse ground truth, thus inflating its F1 score. In contrast, B@ $2\mu$

	Hot-diverse 11%		Hot-narrow 15%		Rare-narrow 73%	
	Hum	Uni	Hum	Uni	Hum	Uni
One2One	93.8	90.4	80.0	61.2	96.0	61.4
PUSL w/o PT	<b>98.0</b>	<b>93.9</b>	<b>93.7</b>	<b>80.1</b>	<b>96.1</b>	<b>75.6</b>

Table 2: Human evaluation of top-5 keyphrase generation on Ads-XMC. Hum: Percentage of human-judged relevant keyphrases. Uni: Percentage of keyphrases present in the label universe. One2One performs poorly in Hot-narrow, an inherently narrow-label subset, where forced overgeneration leads to low relevance and universe coverage.

provides a more balanced evaluation, revealing GROOV’s limitations with the lowest score (24.34) when required to generate more keyphrases.

## 6.2 Early-Termination Bias

We compare PUSL with state-of-the-art GROOV and a plain One2Seq model, all from the One2Seq family. PUSL uniquely generates a flexible range of keyphrases on demand, overcoming self-terminating limitations.

In Table 1, when tasked with generating  $\mu$  or  $2\mu$  keyphrases, PUSL consistently generates more unique keyphrases ( $\#K@_\mu$  and  $\#K@2\mu$ ) across all datasets. In Ads-XMC, where the average keyphrase number is 3, PUSL and the post-trained model generates 3-4 keyphrases when asked for 6 ( $2\mu$ ), while GROOV produces only 1. PUSL’s superior quality is reflected in higher  $F1@O$  and BudgetAccuracy scores. For example, in EURLex4.3k’s Test-Diverse subset, PUSL achieves a  $B@_\mu$  of 88.56, significantly outperforming GROOV (69.06) and One2Seq (77.05). This consistent performance across diverse datasets demonstrates PUSL’s ability to overcome early-termination bias, effectively addressing a critical OXMC limitation when required keyphrase numbers exceed training data averages.

## 6.3 Overgeneration Bias

In evaluations with missing ground truth (Table 1), PUSL outperforms One2One in the high- $\mu$  EURLex4.3k dataset but shows mixed performance in the low- $\mu$  Ads-XMC dataset, highlighting the need for ground truth-free evaluation.

We conduct a human evaluation<sup>3</sup> on Ads-XMC test set samples, assessing keyphrase relevance to input text. We also check if predicted keyphrases exist in a million-level label set (user query universe), as absence may indicate impractical rarity.

<sup>3</sup>We sampled 100 data points for each spilt.

Model	P@O	P@ $\mu$	P@ $2\mu$	B@ $\mu$	B@ $2\mu$	#k@ $\mu$	#k@ $2\mu$
<b>Test-Narrow</b>							
GROOV	<b>54.8</b>	<b>54.8</b>	<b>54.8</b>	18.7	9.4	1.0	1.0
One2Seq	44.7	24.6	24.6	25.5	12.8	1.9	1.9
PUSL	45.7	35.2	17.6	<b>32.4</b>	<b>18.5</b>	<b>2.5</b>	<b>3.4</b>
<b>Test-Diverse</b>							
GROOV	<b>73.4</b>	<b>73.4</b>	<b>73.4</b>	26.0	13.0	1.1	1.1
One2Seq	68.8	41.4	41.4	41.4	20.7	1.9	1.9
PUSL	63.1	72.5	36.3	<b>55.5</b>	<b>33.4</b>	<b>2.5</b>	<b>3.3</b>

Table 3: Comparison of traditional and proposed metrics for OXMC models.

We report the average scores in Table 2. The table shows that One2One scores significantly lower than PUSL in Hot-Narrow human evaluation, likely due to forced generation of irrelevant keyphrases when annotators provide similar keyphrases. For the user query universe evaluation, both models score lower in the Narrow sets, highlighting challenges in predicting diverse keyphrases.

A case study on the “PS4 Shenmue I & II First Limited Edition Sony PlayStation 4 Japan Import” shows that PUSL correctly generated keyphrases like *shenmue ps4*, aligning well with actual user searches, while One2One, despite generating relevant keyphrases, also produced nonexistent ones like *shenmue iii ps4*. Additional case studies and a scalable ground truth-free evaluation are provided in Appendix D.

## 6.4 Evaluation Metrics: P@k vs P@O vs B@k

We compare traditional Precision@k ( $P@k$ ) against our proposed metrics on Ads-XMC (Table 3). The results quantify how  $P@k$  can mislead evaluation with incomplete ground truth - GROOV achieves high  $P@k$  scores (54.8 and 73.4) by generating only one keyphrase (1.0–1.1) that often matches a ground truth label. In contrast,  $B@k$  reveals GROOV’s under-generation through significantly lower scores (9.4 and 13.0).

PUSL generates increasingly more unique keyphrases as  $k$  increases ( $2.5 \rightarrow 3.4$  Test-Narrow,  $2.5 \rightarrow 3.3$  Test-Diverse). While this prolific generation leads to lower  $P@k$  compared to GROOV at higher  $k$  values, PUSL achieves consistently higher  $B@k$  scores (18.5 and 33.4), indicating its predictions better reflect the true label distribution despite incomplete ground truth.

## 6.5 Comparison of PUSL vs GPT-4

In this section, we compare PUSL with zero-shot GPT-4, representing two different paradigms. PUSL is a small domain-specific fine-tuned model,



	GT	Human	PUSL w/o PT
Hot-Diverse	83.0	6.7	3.3
Hot-Narrow	1	1.8	3.1
Rare-Narrow	1	3.3	3.0

Table 4: Comparison of the average number of unique keyphrases for Ground Truth (GT), Human exhaustive annotation, and PUSL (requested to generate 5 keyphrases). PUSL demonstrates consistent performance across scenarios, while humans struggled with Hot-Narrow annotation.

while GPT-4 is a general-purpose LLM using zero-shot learning<sup>4</sup>. This comparison aims to highlight the importance of domain-specific fine-tuning in generating relevant and query-aligned keyphrases.

In the Ads-XMC, GPT-4 achieved perfect human-judgment alignment but generated only 16% of keywords found in the actual user query universe. In contrast, PUSL produced 72% of its predicted keywords within the same universe, demonstrating its effectiveness in generating not only relevant but precisely aligned keyphrases due to domain-specific fine-tuning. This aligns with recent findings on the effectiveness of large language models for keyphrase generation (Song et al., 2023b,a). Although PUSL can be adapted to fine-tune large models, the computational demands are substantial given the OXMC data scale. However, this comparison underlines the benefits of PUSL’s fine-tuning approach in generating keyphrases that meet real-world user needs.

## 6.6 Exhaustive Annotation Experiment

XMC datasets often involve multiple annotators independently selecting data to annotate, leading to potential duplication and missing labels due to self-selection bias. Our proposed PUSL simulates a single annotator exhaustively annotating each data point. We compared this with a human performing the same task. Table 4 reports the average number of unique keyphrases for ground truth (GT, multiple annotators), Human (single exhaustive annotator), and PUSL. Results show that in Hot-Diverse, GT (with thousands of annotators) generates significantly more unique keyphrases than a single Human annotator, who typically experiences cognitive overload after 10 keyphrases, often leading to repetition. For Hot-Narrow, Human produces only 1.8 keyphrases on average, highlighting the challenge of this split.

<sup>4</sup>See Appendix E for the full prompt.

## 7 Discussion

### Annotation Strategies and Self-Selection Bias

OXMC datasets are typically sourced from two main types of data: e-commerce platforms (e.g., amazoncat) and community-driven websites (e.g., Wikipedia articles, Stack Overflow questions). E-commerce data relies on independent user annotations, where each user labels a self-selected subset of items. Community platforms employ collaborative editing, where users iteratively build upon previous annotations. This collaborative process typically yields richer keyphrase sets per item compared to independent annotations, given the same number of annotators. However, OXMC dataset curation obscures this distinction by retaining only unique labels, masking the underlying annotation bias (Marlin et al., 2012; Chen et al., 2023).

### Feedback Loops and Previous Model Bias

Missing labels in e-commerce data also stem from a self-reinforcing cycle between system exposure and user behavior. Search algorithms favor frequently-queried items, increasing their visibility. Enhanced visibility leads to more user interactions, which in turn strengthen these items’ positions in search results. This creates a Matthew effect, where popular items accumulate annotations while less popular ones remain under-annotated. Traditional OXMC datasets only record the final item-label associations without the underlying interaction patterns, perpetuating these biases. Models trained on such biased data may further amplify the disparity through “previous model bias” (Liu et al., 2020), creating feedback loops that reinforce already-popular items. This can lead to “echo chambers” where popular items dominate the label space while niche items become increasingly invisible (Jiang et al., 2019; Ge et al., 2020).

## 8 Conclusion

This paper addresses missing labels in extreme multi-label classification datasets, attributing the issue to self-selection bias. We recast OXMC as Positive Unlabeled Sequence Generation to foster prolific and accurate label prediction, mitigating model laziness caused by missing training labels. Additionally, we propose  $B@k$  metrics for faithful model assessment given incomplete ground-truth labels. Our approach bridges the gap between training label sparsity and inference-time prediction requirements in real-world OXMC applications.

## Limitation

First, we did not conduct missing label analysis on the popular datasets, because the annotation frequency information is unavailable. Second, our current formulation focuses on generating keyphrases for individual items and does not directly address the task of item-item recommendation, because the model’s open-vocabulary generation nature may make up nonexistent items.

## References

- Jessa Bekker and Jesse Davis. 2020. Learning from positive and unlabeled data: A survey. *Machine Learning*, 109(4):719–760.
- Jessa Bekker, Pieter Robberechts, and Jesse Davis. 2019. Beyond the selected completely at random assumption for learning from positive and unlabeled data. In *Joint European conference on machine learning and knowledge discovery in databases*, pages 71–85. Springer.
- Kush Bhatia, Himanshu Jain, Purushottam Kar, Manik Varma, and Prateek Jain. 2015. Sparse local embeddings for extreme multi-label classification. *Advances in neural information processing systems*, 28.
- Charles E. Brown. 1998. *Coefficient of Variation*, pages 155–157. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Ilias Chalkidis, Emmanouil Fergadiotis, Prodromos Malakasiotis, and Ion Androutsopoulos. 2019. [Large-scale multi-label text classification on EU legislation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6314–6322, Florence, Italy. Association for Computational Linguistics.
- Ilias Chalkidis, Manos Fergadiotis, and Ion Androutsopoulos. 2021. [MultiEURLEX - a multi-lingual and multi-label legal document classification dataset for zero-shot cross-lingual transfer](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6974–6996, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Hou Pong Chan, Wang Chen, Lu Wang, and Irwin King. 2019. Neural keyphrase generation via reinforcement learning with adaptive rewards. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2163–2174.
- Wei-Cheng Chang, Daniel Jiang, Hsiang-Fu Yu, Choon Hui Teo, Jiong Zhang, Kai Zhong, Kedarnath Kolluri, Qie Hu, Nikhil Shandilya, Vyacheslav Ievgrafov, et al. 2021. Extreme multi-label learning for semantic matching in product search. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pages 2643–2651.
- Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. 2023. Bias and de-bias in recommender system: A survey and future directions. *ACM Transactions on Information Systems*, 41(3):1–39.
- 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 *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2846–2856, Minneapolis, Minnesota. Association for Computational Linguistics.
- Wang Chen, Yifan Gao, Jiani Zhang, Irwin King, and Michael R. Lyu. 2019b. [Title-guided encoding for keyphrase generation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):6268–6275.
- Ofer Dekel and Ohad Shamir. 2010. [Multiclass-multilabel classification with more classes than examples](#). In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 137–144, Chia Laguna Resort, Sardinia, Italy. PMLR.
- Charles Elkan and Keith Noto. 2008. Learning classifiers from only positive and unlabeled data. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 213–220.
- Yingqiang Ge, Shuya Zhao, Honglu Zhou, Changhua Pei, Fei Sun, Wenwu Ou, and Yongfeng Zhang. 2020. Understanding echo chambers in e-commerce recommender systems. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, pages 2261–2270.
- Nilesh Gupta, Sakina Bohra, Yashoteja Prabhu, Saurabh Purohit, and Manik Varma. 2021. Generalized zero-shot extreme multi-label learning. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 527–535.
- Himanshu Jain, Venkatesh Balasubramanian, Bhanu Chunduri, and Manik Varma. 2019. Slice: Scalable linear extreme classifiers trained on 100 million labels for related searches. In *Proceedings of the twelfth ACM international conference on web search and data mining*, pages 528–536.
- Ray Jiang, Silvia Chiappa, Tor Lattimore, András György, and Pushmeet Kohli. 2019. Degenerate feedback loops in recommender systems. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 383–390.
- Taehee Jung, Joo-kyung Kim, Sungjin Lee, and Dongyeop Kang. 2023. [Cluster-guided label generation in extreme multi-label classification](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*,

- pages 1670–1685, Dubrovnik, Croatia. Association for Computational Linguistics.
- Junmo Kang, Jeonghwan Kim, Suwon Shin, and Sung-Hyon Myaeng. 2021. Leveraging order-free tag relations for context-aware recommendation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3464–3476.
- Ryuichi Kiryo, Gang Niu, Marthinus C Du Plessis, and Masashi Sugiyama. 2017. Positive-unlabeled learning with non-negative risk estimator. *Advances in neural information processing systems*, 30.
- Wee Sun Lee and Bing Liu. 2003. Learning with positive and unlabeled examples using weighted logistic regression. In *ICML, 2003*, pages 448–455.
- Yanfei Lei, Chunming Hu, Guanghui Ma, and Richong Zhang. 2021. Keyphrase extraction with incomplete annotated training data. In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, pages 26–34.
- Bing Liu, Yang Dai, Xiaoli Li, Wee Sun Lee, and Philip S Yu. 2003. Building text classifiers using positive and unlabeled examples. In *Third IEEE international conference on data mining*, pages 179–186. IEEE.
- Dugang Liu, Pengxiang Cheng, Zhenhua Dong, Xiquiang He, Weike Pan, and Zhong Ming. 2020. A general knowledge distillation framework for counterfactual recommendation via uniform data. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, pages 831–840.
- Benjamin Marlin, Richard S Zemel, Sam Roweis, and Malcolm Slaney. 2012. Collaborative filtering and the missing at random assumption. *arXiv preprint arXiv:1206.5267*.
- Julian McAuley and Jure Leskovec. 2013. Hidden factors and hidden topics: understanding rating dimensions with review text. In *Proceedings of the 7th ACM conference on Recommender systems*, pages 165–172.
- Rui Meng, Sanqiang Zhao, Shuguang Han, Daqing He, Peter Brusilovsky, and Yu Chi. 2017. [Deep keyphrase generation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 582–592, Vancouver, Canada. Association for Computational Linguistics.
- Ashirbad Mishra, Soumik Dey, Jinyu Zhao, Hansi Wu, Binbin Li, and Kamesh Madduri. 2024. [Graphite: A graph-based extreme multi-label short text classifier for keyphrase recommendation](#). In *ECAI, Frontiers in Artificial Intelligence and Applications*, page To appear. IOS Press.
- Yashoteja Prabhu, Anil Kag, Shilpa Gopinath, Kunal Dahiya, Shrutendra Harsola, Rahul Agrawal, and Manik Varma. 2018. Extreme multi-label learning with label features for warm-start tagging, ranking & recommendation. In *Proceedings of the eleventh ACM international conference on web search and data mining*, pages 441–449.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Donald B Rubin. 1976. Inference and missing data. *Biometrika*, 63(3):581–592.
- Erik Schultheis and Rohit Babbar. 2021. Unbiased loss functions for multilabel classification with missing labels. *arXiv preprint arXiv:2109.11282*.
- Erik Schultheis, Wojciech Kotłowski, Marek Wydmuch, Rohit Babbar, Strom Borman, and Krzysztof Dembczynski. 2023. Consistent algorithms for multilabel classification with macro-at- $k$  metrics. In *The Twelfth International Conference on Learning Representations*.
- Erik Schultheis, Marek Wydmuch, Rohit Babbar, and Krzysztof Dembczynski. 2022. On missing labels, long-tails and propensities in extreme multi-label classification. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1547–1557.
- Hinrich Schütze, Christopher D Manning, and Prabhakar Raghavan. 2008. *Introduction to information retrieval*, volume 39. Cambridge University Press Cambridge.
- Daniel Simig, Fabio Petroni, Pouya Yanki, Kashyap Popat, Christina Du, Sebastian Riedel, and Majid Yazdani. 2022. [Open vocabulary extreme classification using generative models](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1561–1583, Dublin, Ireland. Association for Computational Linguistics.
- Mingyang Song, Xuelian Geng, Songfang Yao, Shilong Lu, Yi Feng, and Liping Jing. 2023a. Large language models as zero-shot keyphrase extractor: A preliminary empirical study. *arXiv preprint arXiv:2312.15156*.
- Mingyang Song, Haiyun Jiang, Shuming Shi, Songfang Yao, Shilong Lu, Yi Feng, Huafeng Liu, and Liping Jing. 2023b. Is chatgpt a good keyphrase generator? a preliminary study. *arXiv preprint arXiv:2303.13001*.
- Zhenqiao Song, Jiaze Chen, Hao Zhou, and Lei Li. 2021. Triangular bidword generation for sponsored search auction. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pages 707–715.

- Lucas Sterckx, Cornelia Caragea, Thomas Demeester, and Chris Develder. 2016. Supervised keyphrase extraction as positive unlabeled learning. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1924–1929.
- Yue Wang, Jing Li, Hou Pong Chan, Irwin King, Michael R Lyu, and Shuming Shi. 2019. Topic-aware neural keyphrase generation for social media language. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2516–2526.
- Tong Wei and Yu-Feng Li. 2019. Does tail label help for large-scale multi-label learning? *IEEE transactions on neural networks and learning systems*, 31(7):2315–2324.
- Yixing Xu, Chang Xu, Chao Xu, and Dacheng Tao. 2017. Multi-positive and unlabeled learning. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 3182–3188.
- Tiezheng Yu, Hanchao Yu, Davis Liang, Yuning Mao, Shaoliang Nie, Po-Yao Huang, Madian Khabsa, Pascale Fung, and Yi-Chia Wang. 2023. Generating hashtags for short-form videos with guided signals. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9482–9495.
- 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 *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7961–7975, Online. Association for Computational Linguistics.
- Jin Zhang, Jilong Zhang, and Guoqing Chen. 2023. A semantic transfer approach to keyword suggestion for search engine advertising. *Electronic Commerce Research*, pages 1–27.
- 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.
- Hao Zhou, Minlie Huang, Yishun Mao, Changlei Zhu, Peng Shu, and Xiaoyan Zhu. 2019. Domain-constrained advertising keyword generation. In *The World Wide Web Conference*, pages 2448–2459.



---

**Algorithm 2** Dataset Split Procedure

---

**Require:**  $\mathcal{D}_{\text{train}}, \mathcal{D}_{\text{dev}}, \mathcal{D}_{\text{test}}$ 

```
1:  $\triangleright$  Merge and shuffle datasets  $\triangleleft$ 
2:  $\mathcal{D} \leftarrow \text{shuffle}(\mathcal{D}_{\text{train}} \cup \mathcal{D}_{\text{dev}} \cup \mathcal{D}_{\text{test}})$   $\triangleleft$ 
3:  $\triangleright$  Group by text input and concatenate labels  $\triangleleft$ 
4:  $\mathcal{D} \leftarrow \text{GroupBy}(\mathcal{D}, x).Aggregate(\text{concatenate})$   $\triangleleft$ 
5:  $\triangleright$  Uniformly split the dataset  $\triangleleft$ 
6:  $[\mathcal{D}_{\text{train}}, \mathcal{D}_{\text{dev}}, \mathcal{D}_{\text{test}}] \leftarrow \text{UniformSplit}(\mathcal{D})$   $\triangleleft$ 
7:  $\triangleright$  Split based on label number  $\triangleleft$ 
8:  $l \leftarrow \min(\{|y| \mid \forall y \in \mathcal{D}\})$ 
9:  $r \leftarrow \max(\{|y| \mid \forall y \in \mathcal{D}\})$ 
10: for  $k = l$  to  $r$  do
11:    $\mathcal{D}_{\text{train}}, \mathcal{D}_{\text{test}}^k \leftarrow \text{SplitByNum}(\mathcal{D}_{\text{train}}, k)$ 
12: return  $\mathcal{D}_{\text{train}}, \mathcal{D}_{\text{dev}}, \mathcal{D}_{\text{test}}, \{\mathcal{D}_{\text{test}}^k\}_{k=l}^r$ 
```

---

## A Data Details

Prior to (Gupta et al., 2021), XMC datasets were not split from an independent and identically distributed (IID) perspective. Instead, they were split in a way that ensured every label had at least one training point. (Simig et al., 2022) further extended XMC to an open-vocabulary setting where models may be expected to generate unseen labels in the form of multi-token keyphrases. The baselines in this study are all capable of open-vocabulary generalization. Additionally, to evaluate models in a top- $k$  setting under different keyphrase diversity, we merge and re-split the datasets using Algorithm 2.

We also create for each dataset, 2 classes of sub-datasets for reporting aggregated top- $k$  numbers — narrow ( $D_{\text{test}}^n$ ), and diverse ( $D_{\text{test}}^d$ ) which are aggregated from the sub-datasets generated in line 11 of Algorithm 2. In terms of notation:

$$D_{\text{test}}^p = \bigcup_i^j D_{\text{test}}^k \quad (4)$$

where  $i$  and  $j$  denote the lower and upper bounds for the percentile limits for the number of key-phrases per row in the test dataset and  $p \in \{n, d\}$ . For rare ( $D_{\text{test}}^n$ ),  $i = 0$  and  $j = 2 * \mu$  and for diverse ( $D_{\text{test}}^d$ ),  $i = 2 * \mu + 1$  and  $j = 100^{th}$  percentile. The data statistics is in Table 5.

## B Implementation Details

All experiments were conducted using  $4 \times$  Tesla V100-SXM2-32GB GPUs. The PUSL training times for the datasets were as follows:

- Eurlex4.3k: 3 hours
- Wiki: 1 day 20 hours
- AmazonCat: 2 days 18 hours
- Ads-XMC: 1 day 2 hours

We used the T5-base architecture for all compared models with identical hyperparameters. Following (Simig et al., 2022), we set the learning rate to  $2e-4$ , number of epochs to 10, and total batch size to 32. For post-training, we constructed a dataset of 50,000 samples, halved the initial learning rate with linear decay to 0, and trained for 3 epochs. When evaluating top- $k$  models under  $F@O$ , we truncate the predicted keyphrases to a maximum of 100 to avoid memory limitations on our machine.

## C Ablation study on post training

We conducted an ablation study to examine the individual contributions of two data sources, raw and synthetic, used during post-training. This experiment was conducted on a super-diverse test set from Ads-XMC, where the average number of keyphrases per instance is 89.89, compared to 16.91 in the original diverse set. The larger number of keyphrases per instance provides a more comprehensive ground truth, allowing for a more thorough evaluation of the impact of each data source on model performance.

The results, as shown in Table 6, indicate that both data sources contribute to the overall performance, but in distinct ways. The raw data from the original Train-Diverse set helps maintain stable performance on well-represented inputs, though it does not significantly increase the number of generated keyphrases. In contrast, the synthetic data, augmented through rejection sampling, enhances the model’s diversity by increasing the number of keyphrases generated, thus improving performance across various metrics.

When both raw and synthetic data are combined, the model achieves stronger overall performance compared to using either data source individually. This complementary effect suggests that raw data provides stability for common inputs, while synthetic data improves diversity and coverage. It is worth noting that as the model generates more keyphrases, the  $F1@O$  score decreases significantly, reflecting a tendency to favor more prolific models. However, as the model becomes more prolific, the budget accuracy score increases steadily, providing a more faithful and reliable measure of model performance.

To investigate the effects of post-training on the baselines, we conducted additional ablation study for One2Seq, using both raw data from the original Train-Diverse set and synthetic data generated

Dataset	$ \mathcal{D}_{\text{train}} $	$\overline{L}_{\text{train}}$	$ \mathcal{D}_{\text{dev}} $	$\overline{L}_{\text{dev}}$	$ \mathcal{D}_{\text{test}}^{\text{n}} $	$\overline{L}_{\text{test}}^{\text{n}}$	$ \mathcal{D}_{\text{test}}^{\text{d}} $	$\overline{L}_{\text{test}}^{\text{d}}$	$L$
AmazonCat	1,064,452	5.06	25,000	5.07	10,000	5.5	4,313	14.52	13,330
Wikipedia-1M	2,157,967	4.90	142,645	3.95	103,438	3.89	10,128	15.37	960,106
EURLex-4.3K	42,771	15.94	6,000	15.92	3,053	14.04	176	37.53	4,271
Ads-XMC	422,617	3.09	25,000	3.10	5,152	3.15	1,897	16.91	233,612

Table 5:  $|\mathcal{D}_{\text{train}}|$ : number of training instances,  $\overline{L}_{\text{train}}$ : average number of labels per training instance,  $|\mathcal{D}_{\text{dev}}|$ : number of development instances,  $\overline{L}_{\text{dev}}$ : average number of labels per development instance,  $|\mathcal{D}_{\text{test}}^{\text{n}}|$ ,  $|\mathcal{D}_{\text{test}}^{\text{d}}|$ : number of instances in the narrow and diverse test sets, respectively,  $\overline{L}_{\text{test}}^{\text{n}}$ ,  $\overline{L}_{\text{test}}^{\text{d}}$ : average number of labels per instance in the narrow and diverse test sets, respectively,  $L$ : number of unique labels in the whole corpus.

Model	#K@3	#K@6	B@3	B@6	F1@O
PUSL	2.46	3.29	40.32	23.32	<b>35.21</b>
PUSL + raw	2.26	2.75	62.33	36.17	13.94
PUSL + synt	<b>2.82</b>	4.06	77.33	50.67	17.18
PUSL + raw + synt	2.75	<b>4.21</b>	<b>78.67</b>	<b>53.50</b>	18.19

Table 6: Performance comparison of PUSL models with various post-training data.

Model	#K@3	#K@6	B@3	B@6	F1@O
PUSL	2.46	3.29	40.32	23.32	35.21
One2Seq	2.10	2.10	54.33	27.17	10.99
One2Seq + synt	1.77	1.77	48.67	24.33	9.31
One2Seq + raw	1.87	1.87	50.67	25.33	10.18

Table 7: Post-training effects on One2Seq.

through PUSL’s rejection sampling method. However, it is important to note that post-training is not applicable to One2One, as it inherently generates the requested number of keyphrases without flexibility for further augmentation.

As shown in Table 7, simply adding augmented data (either raw or synthetic) to the baseline One2Seq model does not resolve the core challenges of OXMC tasks. In fact, post-training with both data sources resulted in decreased performance for One2Seq, as evidenced by lower metrics across the table. Notably, the number of unique keyphrases generated by One2Seq after post-training (#K) was lower than its original output, indicating a higher repetition rate. We hypothesize that this drop in performance is due to the early-termination bias ingrained in One2Seq during its initial training. Since One2Seq uses an EOS token to conclude keyphrase generation, it tends to terminate prematurely after learning this behavior over a large dataset. During post-training, the smaller amount of augmented data appears insufficient for the model to unlearn this termination behavior, thus limiting its ability to generate more diverse keyphrases



Figure 5: Case study on One2One passing the Uni (user query universe) evaluation but failing on Human evaluation. The item title is “Joystick Rocker Cap Buttons Cover Thumb Stick Grip Cap for PS5 DualSense Edge,” but One2One predicts it as “thumb grip.” While the human annotator considers “thumb grip” as a made-up term, buyers may actually use it to search for joystick caps. This highlights the complexity of evaluating search term predictions, as informal or colloquial terms might be practically useful despite not being formally correct.

## D LLM-as-a-Judge Evaluation

While human evaluation, as conducted in Section 6.3, provides valuable insights, it is time-consuming and challenging to scale. Given the absence of comprehensive labels in the test data, we developed a ground truth-free evaluation method. We leveraged GPT-4 to simulate human evaluation, offering a scalable alternative. Our approach involves prompting GPT-4 with the input text and the keyphrases predicted by the models. GPT-4 then assesses the relevance of each keyphrase to the input, which we denote as binary score  $G$ . This GPT-4-based evaluator has been benchmarked against positive buyer judgment and achieved more than 90% alignment, see (Mishra et al., 2024). We define the final relevance score for each keyphrase

as:

$$Rel = \frac{1}{N} \sum_{i=1}^M \sum_{j=1}^{K_i} (G_{ij} \wedge Uni_{ij}) \quad (5)$$

where  $Rel$  is the overall relevance score,  $M$  is the number of input texts,  $K_i$  is the number of keyphrases for the  $i$ -th input text,  $N = \sum_{i=1}^M K_i$  is the total number of keyphrases across all inputs,  $G_{ij} \in \{0, 1\}$  indicates whether GPT-4 judged the  $j$ -th keyphrase of the  $i$ -th input as relevant, and  $Uni_{ij} \in \{0, 1\}$  indicates whether the same keyphrase is present in the domain-specific universe of queries. The logical AND operation ( $\wedge$ ) ensures that a keyphrase contributes to the relevance score only if it satisfies both general textual relevance and domain-specific considerations.

This methodology allows us to balance general textual relevance ( $G$ ) with domain-specific considerations ( $Uni$ ). While  $G$  captures broader contextual relevance, it may not fully align with domain-specific user preferences, which are better represented by  $Uni$ . The results in Table 8 reveal that the GPT-4 evaluation tends to be more lenient than human assessment, often considering predicted keyphrases as relevant. However, the combined  $Rel$  score provides a more faithful and scalable evaluation metric, effectively incorporating both GPT-4’s textual relevance judgment and domain-specific considerations.

Figure 5 illustrates a compelling case study that highlights the nuances and challenges in evaluating keyphrase predictions. The item in question, “Joy-stick Rocker Cap Buttons Cover Thumb Stick Grip Cap for PS5 DualSense Edge,” was predicted by the One2One model as “thumb grip.” While this prediction passed the Uni (user query universe) evaluation, indicating its presence in domain-specific queries, it failed the human evaluation. The human annotator likely considered “thumb grip” as a non-standard or colloquial term. The combined  $Rel$  score, therefore, provides a more comprehensive assessment that balances linguistic accuracy, domain specificity, and user behavior, offering a more nuanced and practically useful evaluation metric.

## E Prompts for GPT-4

We use the GPT-4 API version gpt-4-0613 for all the experiments.

**GPT-4 as a keyphrase relevance evaluator** Below is the prompt for GPT-4 as a keyphrase relevance evaluator, with templated values to be re-

	Hot-diverse 11%			Hot-narrow 15%			Rare-narrow 73%		
	G	Uni	Rel	G	Uni	Rel	G	Uni	Rel
One2One	<b>92.2</b>	90.4	84.0	90.4	61.2	57.8	93	61.4	57.8
PUSL	90.9	<b>93.9</b>	<b>85.9</b>	<b>92.9</b>	<b>80.1</b>	<b>76.5</b>	<b>93.7</b>	<b>75.6</b>	<b>72.3</b>

Table 8: Scalable evaluation using GPT-4 as a simulated human evaluator, extending Table 2. The combined relevance score ( $Rel$ ) integrates GPT-4’s assessment ( $G$ ) with the domain-specific universe score ( $Uni$ ) using a logical AND operation.

placed by the item title, description, and the model generated keyphrases.

Given an item with title: {title} and description: {desc}, determine whether below keywords are relevant by giving yes or no one by one as answer:

1. {kw1};
2. {kw2};
3. {kw3};
4. {kw4};
5. {kw5}.

###Response:

**GPT-4 as a keyphrase generator** Below is the prompt for GPT-4 as a keyphrase generator, with templated values to be replaced by the item title and description.

Given an item with title: {title} and description: {desc}, give 4 short cpc keywords for the item (all lowercase).  
###Response:

## F Out-of-Time vs Independent-and-identically-distributed Evaluation

Although all models are trained and analyzed under the independent and identically distributed (IID) assumption, they are expected to generalize to future data (out-of-time) that may contain seen/unseen input text, seen/unseen keyphrases, and even completely new entities, such as “*iphone 20*”. Comparing our IID test split and out-of-time split, we found that the performance of PUSL indeed decreases in the future. For instance, in Ads-XMC, the F1 score drops from 32.9 on the Uniform Unseen split to 26.7 on the Future Unseen split (Table 9).

To further investigate this performance drop, we examined the out-of-time Wikipedia-1M dataset (Gupta et al., 2021). We discovered significant issues that led us to exclude it from our main experiment. Firstly, we found that 43% of the input text in the Wikipedia-1M test set have occurred in the training data, but surprisingly, only 0% of the

	Wikipedia-1M			Ads-XMC		
	P	R	F	P	R	F
Uniform Unseen	12.9	8.9	9.7	46.1	34.4	32.9
Future Seen	1.2	1.7	1.3	37.1	44.7	33.4
Future Unseen	11.3	7.3	8.2	21.9	46.9	26.7

Table 9: Evaluation on IID vs future. Seen/Unseen refers to the input text. In uniform split, we group by the text and concatenate the labels, so all text in the Uniform test split are unseen. We report the @4 and @5 scores for Wikipedia-1M and Ads-XMC respectively.

ground truth labels remain the same (Table 9). This stark discrepancy between the input text overlap and the ground truth consistency violates the IID assumption and raises concerns about the validity of the dataset for evaluating model performance on future data. The F1 score on the Wikipedia-1M Future Seen split is only 1.3, compared to 33.4 on Ads-XMC, further highlighting the severity of this issue. In contrast, Ads-XMC OOT exhibits a more reasonable balance, with 53% input text overlap and 45% of the ground-truth labels remaining unchanged. This suggests that Ads-XMC OOT is a more suitable dataset for evaluating the generalization capabilities of XMC models on future data. Secondly, the lack of timestamp information in Wikipedia-1M impedes further analysis of the models’ performance on future data. Without accurate temporal information, it becomes challenging to assess how well the models adapt to evolving trends and new entities over time.

Given these concerns, we decided not to include the Wikipedia-1M OOT dataset in our main experiment. While predicting the future might be an inherent challenge, the inconsistencies and limitations of the Wikipedia-1M OOT dataset make it unsuitable for drawing reliable conclusions about the models’ generalization capabilities. Instead, we focus on evaluating the models using our IID test split and Ads-XMC OOT, which provide a more balanced and representative assessment of their performance.