

REGEN: Zero-Shot Text Classification via Training Data Generation with Progressive Dense Retrieval

Yue Yu¹, Yuchen Zhuang¹, Rongzhi Zhang¹, Yu Meng², Jiaming Shen³, Chao Zhang¹

¹ Georgia Institute of Technology, GA, USA

² University of Illinois at Urbana-Champaign, IL, USA

³ Google Research, NY, USA

{yueyu, yczhuang, rongzhi.zhang, chaozhang}@gatech.edu
yumeng5@illinois.edu, jmshen@google.com

Abstract

With the development of large language models (LLMs), zero-shot learning has attracted much attention for various NLP tasks. Different from prior works that generate training data with billion-scale natural language generation (NLG) models, we propose a retrieval-enhanced framework to create training data from a general-domain unlabeled corpus. To realize this, we first conduct contrastive pretraining to learn an unsupervised dense retriever for extracting most relevant documents using class-descriptive verbalizers. We then further propose two simple strategies, namely *Verbalizer Augmentation with Demonstrations* and *Self-consistency Guided Filtering* to improve the topic coverage of the dataset while removing noisy examples. Experiments on nine datasets demonstrate that REGEN achieves 4.3% gain over strongest baselines and saves around 70% of the time when compared with baselines using large NLG models. Besides, REGEN can be naturally integrated with recently proposed large language models to boost performance¹.

1 Introduction

Text classification serves a fundamental task in Natural Language Processing (NLP) with a broad spectrum of applications. Recently, large pretrained language models (PLMs) (Devlin et al., 2019) have achieved strong performance on text classification with a large amount of task-specific training data. However, in real world scenarios, collecting labeled data can be challenging due to the cost of time, money, and domain expertise.

To reduce the burden of human annotation, we study automatic *dataset generation* for text classification under the zero-shot setting, where no *task-specific* or *cross-task* data is available. Such a setting is different from previous works that use a

large collection of labels from auxiliary tasks for zero-shot text classification (Yin et al., 2019; Gera et al., 2022; Wei et al., 2022; Sanh et al., 2022), and is particularly challenging since we need to adapt the language understanding abilities of PLMs to target classification tasks with minimal supervision.

Prior works on zero-shot synthetic dataset generation mainly fall into two categories: (1) *Generative methods* leverage a billion-scale NLG model to generate class-conditioned texts for PLM fine-tuning (Meng et al., 2022; Ye et al., 2022a,b). While these methods work well on easy tasks (e.g. binary classification), they can be fragile on challenging tasks with more classes, as the generated text can be less discriminative. Besides, the gigantic size of the NLG model will also cause the inefficiency issue. (2) *Mining-based* methods design rule-based regular expressions to extract text from the background corpus as synthesized training data (van de Kar et al., 2022), but these rules are often too simple to capture the complex semantics of text. As a result, the mined dataset contains many incorrectly-labeled data, and the fine-tuned PLM can easily overfit noisy labels.

We design a new framework REGEN² to solve zero-shot text classification. The setting of REGEN is close to the mining-based technique (van de Kar et al., 2022), where a set of class-specific verbalizers and a collection of general-domain unlabeled corpus are available. Motivated by the limitation of hard matching with regular expressions which hardly preserves the meaning of verbalizers, we propose to leverage *dense retrieval* (DR) (Lee et al., 2019; Karpukhin et al., 2020; Xiong et al., 2021; Sun et al., 2022a; Cui et al., 2022), which calculates semantic relevance in a continuous representation space, for dataset curation. With such a *soft matching* mechanism, DR is able to better encode the category-specific semantics and thus fetch the relevant documents from the corpus. To integrate

¹The code and unlabeled corpus will be released in <https://github.com/yueyu1030/ReGen>.

²Retrieval-Enhanced Zero-shot Data Generation.

DR with the target classification task, we employ two PLMs: one retrieval model (R_θ) to extract the most relevant documents from the unlabeled corpus for synthetic dataset curation, and one classification model (C_ϕ) to be fine-tuned on the generated synthetic dataset to perform the downstream task. Before performing text retrieval, we first conduct contrastive learning on the unlabeled corpus to further pretrain the retrieval model R_θ for producing better sequence embeddings. Then, with the retrieval model, we use the verbalizers from each class as queries to retrieve relevant documents from the unlabeled corpus, which will be used as the training data for target tasks.

Simply fine-tuning the classifier on the above training data may yield limited performance, as the verbalizers are often too generic to cover all the category-related topics (*e.g.*, the word ‘sports’ alone does not cover concrete types of sports). Thus, the retrieved data may contain noisy and irrelevant documents. To enhance the quality of the synthetic dataset, we conduct multi-step retrieval with two additional strategies to strengthen our framework: (1) we *augment* the verbalizer with the retrieved documents from the previous round as additional information (Xu and Croft, 2017) to enrich its representation, which allows for extracting more relevant documents for the downstream task. (2) we exploit *self-consistency* to filter the potentially incorrect examples when the pseudo labels produced by the retrieval model (R_θ) and the classifier (C_ϕ) disagree with each other. We note that REGEN *does not* use annotated labels from any other tasks, making it applicable to the true zero-shot learning. Besides, REGEN only requires two BERT_{base} scale PLMs, which is efficient compared with methods using large NLG models.

Our contribution can be summarized as follows: (1) We propose REGEN, a framework for zero-shot dataset generation with a general-domain corpus and retrieval-enhanced language models. (2) We develop two additional techniques, namely verbalizer augmentation with demonstration and self-consistency guided filtering to improve the quality of the synthetic dataset. (3) We evaluate REGEN on *nine* NLP classification tasks to verify its efficacy. We also conduct detailed analysis to justify the role of different components as well as the robustness of REGEN over different verbalizers.

2 Related Work

Zero-shot Text Classification (ZSTC) aims to categorize the text document without using task-specific labeled data. With pretrained language models, a plenty of works attempted to convert the classification task into other formats such as masked language modeling (Hu et al., 2022; Gao et al., 2021a), question answering (Zhong et al., 2021; Wei et al., 2022; Sanh et al., 2022) or entailment (Yin et al., 2019; Gera et al., 2022) for zero-shot learning. These works are orthogonal to REGEN as we do not directly perform inference and do not leverage human annotations from additional tasks.

More relevant to us, there are some recent studies that perform ZSTC via generating a task-specific dataset using NLG models, which is then used to fine-tune a classifier for the target task such as text classification (Ye et al., 2022a,b; Meng et al., 2022), sentence similarity calculation (Schick and Schütze, 2021b), commonsense reasoning (Yang et al., 2020; Kan et al., 2021), and instruction-based tuning (Wang et al., 2022). Unfortunately, the generation step is time-consuming and the quality of the generated text can be less satisfactory in capturing fine-grained semantics. The most relevant work to us is (van de Kar et al., 2022), which also extracts documents from the unlabeled corpus to form the training set. But it simply uses regular expressions to mine documents and cannot fully capture the contextual information of verbalizers. Instead, we leverage dense retrieval for concept understanding and obtain the most relevant documents, which is combined with verbalizer augmentation to improve retrieval quality.

On the other hand, retrieval-augmented language models have been used in language modeling (Khandelwal et al., 2020; Borgeaud et al., 2022), OpenQA (Jiang et al., 2022; Sachan et al., 2021), information extraction (Zhuang et al., 2022) and knowledge-intensive tasks (Lewis et al., 2020; Izacard et al., 2022b), where tokens or documents are retrieved based on contextual representations and are used as additional inputs to support target tasks. While such a paradigm has also been explored for zero-shot learning, it is mainly used for zero-shot prompt-based inference (Shi et al., 2022; Chen et al., 2022). Instead, we empirically demonstrate the efficacy of retrieval-enhanced learning for zero-shot dataset curation with an unsupervised dense retrieval model.

3 Preliminaries

◇ **Setup.** We focus on synthesizing a task-specific dataset for text classification (Meng et al., 2022; van de Kar et al., 2022). Besides, we stick to the *strict* zero-shot setup (Perez et al., 2021), where *no labeled examples* from either target tasks or other tasks are available.

◇ **Available Resources.** Besides annotated labels, the availability of massive task-specific unlabeled data is also a rarity — in prior works, such unlabeled data is obtained via removing the ground-truth label from the original dataset (Meng et al., 2020b), and can be scarce in real zero-shot settings (Tam et al., 2021). The most accessible information is a collection of general-domain unlabeled corpus \mathcal{D} (e.g., WIKI), which is freely available online and has been used for pretraining (Devlin et al., 2019; Gururangan et al., 2020). Recent works have also use such an external corpus for zero-shot learning (Shi et al., 2022; van de Kar et al., 2022).

◇ **Task Formulation.** With the above discussion, we consider the classification task where we are given the label set $\mathcal{Y} = \{1, 2, \dots, c\}$ (c is the number of classes), and a mapping $\mathcal{M} : \mathcal{Y} \rightarrow \mathcal{W}$ that converts each label $y \in \mathcal{Y}$ into a class-descriptive verbalizer $w_y \in \mathcal{W}$. We also assume a general-domain unlabeled corpus \mathcal{D} is available. We seek to curate training data \mathcal{T} from \mathcal{D} and learn a PLM C_ϕ which will be fine-tuned as the classifier.

◇ **Backgrounds for Dense Retrieval (DR).** In dense retrieval (Lee et al., 2019), the PLM is used to represent queries and documents in dense vectors. The relevance score $f(q, d)$ is calculated with a scoring function (e.g., dot product) between query and document vectors

$$f(q, d) = \text{sim}(R_\theta(q), R_\theta(d)), \quad (1)$$

where the embedding of the [CLS] token from the final layer of R_θ is used as the representation for both queries and documents. In practice, the documents are encoded offline, and can be efficiently retrieved using approximate nearest neighbor search (ANN) with the queries (Johnson et al., 2021).

4 Method

In this section, we present REGEN (our framework) and introduce the major components.

4.1 Contrastive Pretraining for Retriever R_θ

Directly using BERT for retrieval can lead to unsatisfactory results since BERT embeddings are

not tailored for retrieval application (Gao et al., 2021b). To effectively train a dense retrieval model *without relevance supervision*, we hypothesize that two sentences from the same document share similar semantics as they may describe the same topic. Then, we continuously pretrain the PLM on the corpus \mathcal{D} with contrastive learning (Gao and Callan, 2022; Izacard et al., 2022a; Yu et al., 2022b): Given a document $d_i \in \mathcal{D}$, the positive pair (x_i, x_i^+) is constructed by randomly sampling two disjoint sentences from d_i . Let $\mathbf{h}_i = R_\theta(x_i)$, $\mathbf{h}_i^+ = R_\theta(x_i^+)$ denote the representation of x_i and x_i^+ encoded by the retriever R_θ , the training objective of contrastive learning for pair (x_i, x_i^+) with a mini-batch of N pairs is:

$$\ell_{\text{cl}} = -\log \frac{e^{\text{sim}(\mathbf{h}_i, \mathbf{h}_i^+)/\tau}}{\sum_{j=1}^N e^{\text{sim}(\mathbf{h}_i, \mathbf{h}_j^+)/\tau}}, \quad (2)$$

where we use in-batch instances as negative samples (Gillick et al., 2019), $\text{sim}(\mathbf{h}_i, \mathbf{h}_i^+) = \mathbf{h}_i^\top \mathbf{h}_i^+$ is the dot product, and $\tau = 1$ is the parameter for temperature. Contrastive learning improves the representations by promoting the alignment of similar text sequences and the uniformity of unrelated text sequences, thus enhancing the embedding quality for documents in \mathcal{D} .

4.2 Overall Pipeline

With a pretrained retrieval model R_θ , REGEN follows a *retrieve-then-finetune* pipeline to curate the training data from the corpus \mathcal{D} which will be used to finetune the PLM classifier C_ϕ . The details of our framework are described as follows.

Document Retrieval with Verbalizers. With the class-specific verbalizers, we construct the input queries for each class to retrieve the relevant documents from \mathcal{D} . Formally, the query for the i -th class ($1 \leq i \leq c$) can be expressed as

$$q_i = [\text{CLS}] \circ \mathcal{P}(w_i) \circ [\text{SEP}],$$

where $\mathcal{P}(w_i)$ is the template for the corresponding class with the verbalizer w_i and \circ stands for the concatenation operation. For instance, a query for the binary sentiment classification can be formulated as $q_i = [\text{CLS}] \text{ It was } w_i [\text{SEP}]$, where w_1 and w_2 ($c = 2$ in this case) stand for the verbalizers, namely “*bad*” (negative) and “*great*” (positive), respectively. By feeding the class-dependent query into the retriever R_θ , we expect the retriever to understand its contextualized semantics (Rubin et al., 2022), and extract the relevant documents from the corpus which serve as training examples for the cor-

Algorithm 1: Process of REGEN.

Input: \mathcal{D} : Unlabeled Corpus; \mathcal{Y} : Label space; \mathcal{P} : Verbalizers; R_θ : Retrieval Model; C_ϕ : Classification Model; T : Rounds of Retrieval.

// Step 0: Contrastive Learning.

Pretrain R_θ with Contrastive Learning via Eq. 2.

for $t = 1, 2, \dots, T$ **do**

// Step 1: (Multi-step) Document Retrieval.

if $t = 1$ **then**

 Retrieve Documents \mathcal{T}^1 with \mathcal{P} via Eq. 3.

else

 Retrieve Documents \mathcal{T}^t with \mathcal{P} and $\tilde{\mathcal{T}}^{t-1}$ via Eq.6. *// Verbalizer Augmentation.*

// Step 2: Document Filtering.

 Obtain Filtered Dataset $\tilde{\mathcal{T}}^t$ via Eq. 7.

// Step 3: Language Model Fine-tuning.

 Fine-tune PLM C_ϕ^t with $\tilde{\mathcal{T}}^t$ via Eq. 4.

Output: The dataset $\tilde{\mathcal{T}}^t$ and the PLM classifier C_ϕ^t .

responding category. For the i -th class, the initial retrieved dataset $\mathcal{T}_i^1 \subset \mathcal{D}$ can be written as

$$\mathcal{T}_i^1 = \text{Top-k } f(q_i, d), \quad (3)$$

where $f(q, d)$ is defined in Eq. 1. The full retrieved dataset can be expressed as $\mathcal{T}^1 = \bigcup_{1 \leq i \leq c} \mathcal{T}_i^1$.

Fine-Tuning PLM with Curated Data. After obtaining the training data \mathcal{T} from the corpus³, one can fine-tune a PLM classifier C_ϕ for the downstream task. To achieve better fine-tuning stability and generalization, we adopt the simple *label smoothing* (LS) technique (Müller et al., 2019), which mixes the one-hot labels with uniform vectors. For a training example $(x, y) \in \mathcal{T}$, C_ϕ is trained to minimize the divergence between the label and the classifier’s prediction $p_\phi(x)$ as

$$\min_{\phi} \ell_{\text{ft}} = - \sum_{j=1}^c q_j \log(p_\phi(x)_j), \quad (4)$$

where $q_j = \mathbb{1}(j = y)(1 - \alpha) + \alpha/c$ is the smoothed label and $\alpha = 0.1$ is the smoothing term. LS prevents C_ϕ from overfitting to training data by forcing it to produce less confident predictions.

4.3 Progressive Training Data Curation via Multi-step Dense Retrieval

Although the aforementioned pipeline can retrieve a set of documents used for training (\mathcal{T}^1), the performance can still be suboptimal because (1) the training set only have *limited coverage* as the verbalizers only contains few key words which is too specific to fully represent the categorical informa-

³Here we omit the superscript for \mathcal{T} as the fine-tuning procedure remains the same for all rounds and generated datasets.

tion. (2) the training set still contain *noisy* or *task-irrelevant* documents as the R_θ may not always retrieve texts pertaining to the desired class. To overcome these drawbacks, we perform document retrieval for multiple rounds, employing two additional strategies as described below.

Verbalizer Augmentation with Demonstrations.

The verbalizers often contain only a few words and are insufficient to perfectly reflect the underlying information. Motivated by the recently proposed demonstration-based learning (Brown et al., 2020; Min et al., 2022) which augments the input with labeled examples to support in-context learning, we aim to enrich verbalizers with top retrieved documents for improving their representations (Yu et al., 2021), and thus enhancing the quality of the retrieved data. Specifically, in the t -th ($t > 1$) round, we use the retrieved documents from the $t-1$ round as demonstrations to augment the verbalizer for the i -th class as⁴

$$q_{i,j}^t = [\text{CLS}] \circ \mathcal{P}(w_i) \circ [\text{SEP}] \circ d_{i,j}^{t-1} \circ [\text{SEP}], \quad (5)$$

where $d_{i,j}^{t-1}$ is the j -th documents for the i -th class in the previous dataset $\tilde{\mathcal{T}}^{t-1}$. With the augmented queries, \mathcal{T}_i^t and \mathcal{T}^t are obtained via combining the retrieved documents as

$$\mathcal{T}_i^t = \bigcup_j (\text{Top-k } f(q_{i,j}^t, d)), \mathcal{T}^t = \bigcup_{1 \leq i \leq c} \mathcal{T}_i^t. \quad (6)$$

Filtering Noisy Data guided via Self-consistency.

The above retrieval process may also introduce noisy examples due to the limited capability of the retrieval model. While the label smoothing in Eq. 4 can mitigate this issue during fine-tuning, it is a generic technique without considering task-specific knowledge. To further fulfill the denoising purpose, we simply leverage the classifier from the previous round and exploit the *consistency* between the retriever and classifier to identify potential incorrect examples. For the example from the t -th round ($t > 1$) denoted as $(x^t, y^t) \in \mathcal{T}^t$ where y^t is the label for the augmented verbalizer, we generate the predicted label using the classifier C_ϕ^{t-1} from the previous round⁵ as $\hat{y}^{t-1} = \text{argmax } p_\phi^{t-1}(x^t)$.

Then, the filtered dataset $\tilde{\mathcal{T}}^t$ is expressed as

$$\tilde{\mathcal{T}}^t = \{(x^t, y^t) \in \mathcal{T}^t \mid \text{argmax } p_\phi^{t-1}(x^t) = y^t\}. \quad (7)$$

To interpret Eq. 7, we only preserve examples

⁴We obtain *multiple* queries for each class after this step.

⁵When $t = 1$, we use the zero-shot prompting model as the classifier due to the absence of the ‘previous model’.

Dataset	Task	Class	# Test	Metric
AGNews	News Topic	4	7.6k	Accuracy
DBPedia	Wikipedia Topic	14	70k	Accuracy
Yahoo Topics	Web QA Topic	10	60k	Accuracy
NYT	News Topic	9	30k	F1
IMDB	Movie Review Sentiment	2	25k	Accuracy
MR	Movie Review Sentiment	2	2k	Accuracy
SST-2	Movie Review Sentiment	2	0.8k	Accuracy
Amazon	Product Review Sentiment	2	40k	Accuracy
Yelp	Restaurant Review Sentiment	2	38k	Accuracy

Table 1: Dataset statistics.

where the prediction from the previous classifier \hat{y}^{t-1} and the retrieved label y^t are *consistent* to fine-tune the classifier C_ϕ , thus serving as an additional protection for C_ϕ against overfitting to label noises.

4.4 Overall Algorithm

The procedure of REGEN is summarized in Algorithm 1. Note that the retrieval model pretraining and corpus indexing only need to be done *once* before applying to all datasets. In each round of retrieval, it only needs one extra ANN retrieval operation per query, which is efficiently supported by FAISS (Johnson et al., 2021). We conduct the efficiency study in the Section 5.9.

5 Experiments

5.1 Experimental Setups

◇ **Datasets.** In this work, we select **AG News** (Zhang et al., 2015), **DBPedia** (Lehmann et al., 2015), **Yahoo** (Zhang et al., 2015) and **NYT** (Meng et al., 2020a) for topic classification, and **IMDB** (Maas et al., 2011), **SST-2** (Socher et al., 2013), **Amazon** (McAuley and Leskovec, 2013)⁶, **MR** (Pang and Lee, 2005), **Yelp** (Zhang et al., 2015) for sentiment analysis. All the datasets are in English. We report performance on the test set when available, falling back to the validation set for SST-2. The details for these datasets can be found in table 1.

◇ **Corpus.** We follow (Shi et al., 2022; van de Kar et al., 2022) to obtain a heterogeneous collection of text that are broadly relevant to tasks in our experiments as the general-domain unlabeled corpus \mathcal{D} . Specifically, we select WIKI (Petroni et al., 2021), subsets of REVIEWS (He and McAuley, 2016) and REALNEWS (Zellers et al., 2019) to form the corpus. The detailed information and preprocessing steps for these corpora are shown in Appendix B.

⁶We follow (Hu et al., 2022) to subsample a 40K subset from the original 400K test data for faster evaluations, which has little effect on the average performance in our pilot studies.

◇ **Metrics.** We use F1 score as the metric for NYT as the label distribution is imbalanced. Accuracy is used for the remaining tasks.

◇ **Baselines.** We consider various baselines, including both zero-shot inference and dataset generation methods. Details of the baselines are in Appendix C. We also list the results with extra resources (*e.g.* large PLMs, task-specific samples, or knowledge bases), but only for reference purposes, *since we do not claim REGEN achieves state-of-the-art performance on zero-shot text classification. Rather, we consider REGEN as a better approach to synthesizing datasets in a zero-shot manner for text classification tasks.*

◇ **Implementation Details.** For implementation, we use PyTorch (Paszke et al., 2019) and HuggingFace (Wolf et al., 2019). We set the retrieval rounds $T = 3$, the k used in ANN in Eq. 3 to 100 for the 1st round and 20 for later rounds in Eq. 6. The number of the training data per class is set to no more than 3000 (Meng et al., 2022). Under the zero-shot learning setting, we keep all hyperparameters the *same* across all tasks due to the lack of validation sets. In principle, REGEN is compatible with any dense retriever R_θ and classifier C_ϕ . In this work, we initialize R_θ from Condenser (Gao and Callan, 2021) and fine-tune RoBERTa-base (Liu et al., 2019) as C_ϕ . See App. D for details.

5.2 Main Experiment Results

The results of REGEN and compared baselines on nine tasks are in Table 2. From these results, we have the following observations:

- (1) REGEN significantly surpasses fair baselines on average of nine datasets, and often achieves comparable or even better results against methods using extra task-specific information. Compared with our direct baseline (van de Kar et al., 2022) using regular expressions to mine training data, REGEN achieves 4.3% gain on average. The gain is more notable (6.8%) for topic classification with more classes. These results justify that dense retrieval serves as a more flexible way to understand the category and can extract training data being *semantically closer* to the target topics.
- (2) SuperGen (Meng et al., 2022) achieves strong results on sentiment tasks. However, its performance diminishes for multi-class topic classification, suggesting that NLG-based dataset generation methods may struggle to produce sufficiently accurate and distinct texts for fine-grained classification.

Task (→)	Topic Classification					Sentiment Classification					All	
Method (↓) / Dataset (→)	AG News	DBPedia	Yahoo	NYT	Avg.	IMDB	MR	SST-2	Amazon	Yelp	Avg.	Avg.
<i>Zero-shot Learning via Direct Inference on Test Data</i>												
NSP-BERT (2022b)	78.1	69.4	47.0	54.6	62.3	73.1	74.4	75.6	69.4	66.3	71.8	67.5
Prompt (2021a)	73.2	71.3	44.1	57.4	61.5	74.8	73.2	75.9	80.2	78.1	76.4	68.9
KNN-Prompt (2022)	78.8	—	51.0	—	—	—	78.2	84.2	85.7	—	—	—
GPT-3 [‡] (2021)	73.9	59.7	54.7	57.0	61.3	75.8	76.3	87.2	75.0	78.5	78.6	69.9
<i>Zero-shot Learning via Generating Task-specific Datasets</i>												
SuperGen [†] (2022)	77.4±1.5	66.5±2.0	40.8±1.5	53.9±1.5	59.7	85.8±1.6	81.9±0.9	88.6±0.5	91.0±0.9	93.6±0.6	88.1	73.9
Mining [*] (2022)	79.2	80.4	56.1	—	—	86.7	80.5	85.6	92.0	92.0	87.3	—
Mining ^{*‡} (Our Relmp.)	79.7±1.0	82.1±0.6	57.0±0.6	68.6±0.9	71.9	87.1±0.6	79.9±0.7	85.0±0.6	92.1±0.5	92.3±0.5	87.2	79.6
REGEN (Our Method)	85.0±0.8	87.6±0.9	59.4±0.8	74.5±1.1	76.6	89.9±0.5	82.5±0.7	88.9±0.4	92.3±0.4	93.0±0.5	89.3	83.0
<i>For Reference Only: Using labeled data from other tasks / task-specific corpus / external knowledge base.</i>												
TE-NLI (Best) [†] (2019)	78.0	73.0	43.8	70.7	66.4	64.6	68.3	68.6	76.7	73.5	70.3	68.6
NLI-ST ^{†‡} (2022)	76.5	92.2	59.8	—	—	92.5	—	—	94.3	—	—	—
KPT ^{‡,§} (2022)	84.8	82.2	61.6	72.1	75.2	91.2	—	—	92.8	—	—	—
LOTClass [‡] (2020b)	86.2	91.1	55.7	49.5	70.7	86.5	70.8	80.9	91.7	87.6	83.5	77.1
X-Class [‡] (2021)	85.7	91.3	50.5	68.5	74.0	89.0	78.8	84.8	90.4	90.0	86.5	80.3

Table 2: Main results. We report average performance and standard deviation across 5 runs *if fine-tuning is applied*. *: concurrent work, ‡: use the same corpus and template as REGEN for *fair comparisons*, †: use labeled data from auxiliary tasks, ‡: use task-specific corpus, ‡: use billion-scale PLMs, §: use additional knowledge base.

Method	AG News	DBPedia	SST-2	Yelp
REGEN	85.0	87.6	88.9	93.0
w/o Data Curation (DC)	70.9	68.8	69.2	75.5
w/o Multi-step Retrieval (MSR)	83.0	83.6	85.9	90.9
w/o Label Smoothing (LS)	84.5	86.1	88.0	91.7

Table 3: Ablation Study. For w/o DC, we use R_θ to calculate similarity between samples and labels for zero-shot learning. For w/o MSR, we only retrieve *the same size of data* as REGEN for one round with verbalizers. For w/o LS, we use one-hot labels for fine-tuning.




Figure 1: Effect of self-consistency guided filtering.

(3) REGEN also delivers competitive performance against zero-shot learning and weakly-supervised text classification baselines without requiring additional resources, such as larger language models or task-specific unlabeled data. This suggests that dataset generation serves as an alternative approach for zero-shot text classification.

5.3 Ablation Studies

Effect of Different Components. Table 3 shows the result of ablation studies on four datasets⁷, which demonstrates the superiority of retrieving texts from the corpus for training data creation as well as conducting multi-step retrieval. Besides, label smoothing also results in performance gain as it mitigates the effect of noisy labels for fine-tuning.

Besides, we plot the result over different rounds

⁷More results on other datasets are in Appendix H.

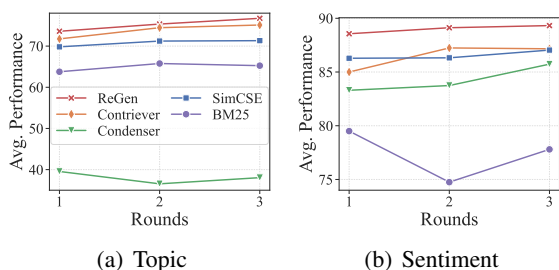


Figure 2: Effect of different dense retrieval models R_θ .

of retrieval in Fig. 1. It is clear that both multi-step retrieval and filtering progressively enhance the performance of target tasks, justifying their necessity for improving the quality of training data. We have also attempted to conduct more retrieval rounds, but do not observe significant performance gains.

Study of Dense Retrievers. We compare the retrieval model R_θ with other off-the-shelf unsupervised retrieval models. Here we choose one sparse model BM25 (Robertson et al., 2004) and three DR models: Condenser (Gao and Callan, 2021), SimCSE (Gao et al., 2021b), and Contriever (Izacard et al., 2022a). From Figure 2, we observe that the performance of BM25 is not satisfactory, since simply using lexical similarity is insufficient to retrieve a set of diverse documents for fine-tuning. Besides, our retrieval model outperforms other unsupervised DR models for two reasons: (1) Condenser and SimCSE are pretrained over *short sentences*, and the learning objective is suboptimal for long documents; (2) these models are not pretrained on the corpus used in our study and suffer from the *distribution shifts* (Yu et al., 2022b). Instead, our strategy can better adapt the PLM for the retrieval task.




Figure 3: Effect of the training data size.




Figure 4: Accuracy on IMDB/AG News fine-tuned on the few labeled samples only vs. on the few-shot and generated dataset with varying amount of labeled data.

In the following sections, we mainly compare REGEN with Mining (van de Kar et al., 2022) and SuperGen (Meng et al., 2022) as they are closest baselines to us.

5.4 Effect of the Amount of Generated Data

Figure 3 shows the results of using different amount of training data (after filtering). Overall, we find that the performance first improves significantly when the number of training data is small (e.g., 100), then becomes stable with more retrieved data. This is because with too many generated data, it may also introduce more label noise and reduce the quality of training data. Nevertheless, REGEN outperforms baselines under different volumes of training samples, justifying its advantage.

5.5 Fusing REGEN with Large Language Models (LLMs)

In this section, we give a simple demonstration of how to leverage recently-proposed large language models (e.g. GPT-4 (OpenAI, 2023)) to further boost the performance. As LLMs have demonstrated strong ability for text generation, we use them to augment the verbalizer before retrieving documents from the general-domain corpus. The details are in Appendix E.3.

From Table 4, we observe that expanded verbalizers lead to consistent performance gains on two datasets. Although the scale of the improvement is not that significant, it shows some effectiveness with such cheap plug-in techniques of using LLMs

Dataset	AG News		DBPedia	
	REGEN	REGEN+LLM	REGEN	REGEN+LLM
Accuracy	85.0 \pm 0.8	85.4 \pm 0.5	87.6 \pm 0.9	88.5 \pm 0.8

Table 4: Effect of using Large Language Models for Verbalizer Expansion.

Dataset	Verbalizer Group	Mining	SuperGen	REGEN
IMDB	# 0 (Original)	87.1	85.8	89.9
	# 2	88.3	82.7	87.9
	# 3	86.0	80.6	90.3
	# 4	89.0	89.1	90.1
	Avg. \pm Std.	87.6 \pm 1.3	84.5 \pm 3.6	89.6\pm1.2
AG News	# 0 (Original)	79.7	77.4	85.0
	# 1	82.5	75.2	82.7
	# 2	83.9	77.8	85.1
	# 3	77.7	72.2	83.6
	Avg. \pm Std.	80.9 \pm 2.7	75.6 \pm 2.5	84.2\pm1.2

Table 5: Results with different verbalizers. The number for each prompt group is the averaged performance across 5 runs; Avg. \pm Std. is calculated over four groups.

for boosting REGEN.

5.6 Using REGEN in Few-Shot Settings

REGEN can also be combined with a few labeled examples to improve the performance. We follow (Meng et al., 2022) to fine-tune C_ϕ with few-shot examples and the synthetic dataset (Details in Appendix E.1) using IMDB and AG News as examples. From Fig. 4, we observe that REGEN improves over the vanilla few-shot fine-tuning under all studied regions (32 to 512 labels per class), while baselines cannot further promote the performance with more training samples. Quantitatively, the performance of REGEN is equivalent to that of fine-tuning with 128-256 labeled documents per class. With 32 labels per class, REGEN achieves comparable performance of vanilla fine-tuning with 4x-8x labeled examples. These results verify that REGEN promotes label efficiency of PLMs.

5.7 Robustness over Different Verbalizers

As REGEN and zero-shot dataset generation methods always rely on a class-dependent verbalizer to steer the whole process, we study the impact of different verbalizers on the final performance. We use IMDB and AG News as two datasets, and create three different groups of verbalizers other than the default ones for comparison (Details in Appendix E.2). From Table 5, we observe that REGEN generally outperforms baselines on 7 out of 8 cases. REGEN also has *lower* performance variance across four groups of verbalizers. These results reveal that REGEN does not rely on specific designs of verbalizers, and are more robust over




Figure 5: (a) Performance of REGEN and Mining using only subset of corpus \mathcal{D} . N/W/R stands for REAL-NEWS/WIKI/REVIEWS, respectively. (b) The relation on the performance gap and lexical similarity between the corpus and target tasks.

Operation	Mining	Supergen	REGEN
Pretraining	—	—	23h
Indexing of Corpus/Per doc	—	—	6h/4ms
Curating Dataset Per Task	1.4h	20.4h	0.6h
Filtering Per Task	0.2h	0.1h	0.5h
Model Fine-tuning Per Task	0.4h	0.3h	0.7h
Total Time (for all Tasks)	10h	104h	38h

Table 6: Efficiency Study. For REGEN, the average time per task of curating dataset, filtering and fine-tuning is accumulated over 3 rounds.

different verbalizers.

5.8 The Effect on General-domain Corpus \mathcal{D}

We study the effect of corpus \mathcal{D} by conducting retrieval on different subsets from \mathcal{D} . As shown in Figure 5(a), we observe better performance when the corpus aligns with the target task well (e.g. NEWS for AG News). This is expected as the model suffers less from the distribution shift issue. Besides, REGEN outperforms the mining method under all settings, justifying its superior ability to retrieve relevant text even if there is a domain mismatch between the task and corpus.

Fig. 5(b) exhibits the relation on the lexical similarity (measured by weighted Jaccard score), and the performance gap between REGEN and fully-supervised BERT (details in Appendix G). Overall, there is a negative correlation among performance gaps and the distribution similarities, as REGEN performs closer to fully-supervised models on tasks where task-specific documents share more similar lexical patterns with the general-domain corpus.

5.9 Efficiency Studies

Table 6 measures the efficiency of REGEN and baselines. While the pretraining and indexing corpus for REGEN can be time-consuming, it only needs to be done once, thus the overall running time of REGEN is significant lower than the baseline using large NLG models (Meng et al., 2022).

Dataset	Metrics	Mining	SuperGen	REGEN
Sentiment	Correctness (\uparrow)	0.815	0.971	0.986
	Diversity (\downarrow)	0.144	0.915	0.361
	Distribution Sim. (\uparrow)	0.856	0.803	0.865
Topic	Correctness (\uparrow)	0.759	0.626	0.860
	Diversity (\downarrow)	0.132	0.767	0.346
	Distribution Sim. (\uparrow)	0.748	0.648	0.757

Table 7: Automatic evaluation results on three metrics.

Dataset	Metrics	Mining	SuperGen	REGEN
Sentiment	Correctness (\uparrow)	1.46	1.95	1.94
	Diversity (\uparrow)	2.00	0.75	2.00
	Informativeness (\uparrow)	1.40	1.90	1.92
AG News	Correctness (\uparrow)	1.78	1.74	1.94
	Diversity (\uparrow)	1.62	0.94	1.88
	Informativeness (\uparrow)	1.63	1.43	1.82

Table 8: Human evaluation results on three metrics. (The full score is 2)

Compare with the mining-based method, although REGEN costs longer time in total, we think it is worthwhile as REGEN outperforms it on all nine tasks studied in this work.

5.10 Quality Analysis of Synthetic Datasets

We provide other measurements to better evaluate the quality of the generated dataset of REGEN and baselines (Ye et al., 2022a).

Automatic Evaluations. We first measure the quality of the dataset from three perspectives: *correctness*, *diversity*, *distribution similarity*. The details are shown in Appendix I.1. Overall, the diversity of generated text from NLG models (Meng et al., 2022) is not satisfactory, and the correctness of text from NLG models is also not guaranteed for topic classification tasks. For the mining-based method, despite it achieves better diversity, the performances on other two metrics are worse. As a result, REGEN surpasses it on these tasks.

Human Evaluations. We also conduct human evaluations to evaluate the quality of the synthetic dataset using AG News and Sentiment datasets as two examples. For each class, we randomly sample 25 documents and ask 4 human volunteers to evaluate the dataset from three perspectives: *Correctness*, *Informativeness* and *Diversity* (details in Appendix I.2). The mean ratings are shown in Table 8. The average Fleiss’ Kappa (Fleiss, 1971) for correctness, informativeness and diversity are 0.53/0.57/0.58 (Moderate Agreement), respectively. Overall, the dataset curated by REGEN has the best informativeness and diversity, while has a competitive result on correctness score. These results indicate that REGEN improves over previous works for curating a better dataset to tackle the downstream

tasks. Detail cases of samples from the synthetic datasets can be found at Appendix J.

6 Discussion and Conclusion

6.1 Discussion

Extending REGEN to Specific Domains. The REGEN framework is versatile and can be applied to various domains beyond our experiments. For example, it is possible to extend REGEN to zero-shot biomedical text classification (Cohan et al., 2020) using the publicly available PubMed articles as the unlabeled corpus.

Verbalizers Selection for REGEN. All the verbalizers used in this work are from the prior works (Hu et al., 2022; Schick and Schütze, 2021a) to circumvent manual prompt engineering and ensure a fair comparison. For those datasets where verbalizers are not given, we can adopt automatic verbalizer and template generation approaches (Gao et al., 2021a) to generate verbalizers for retrieving relevant documents.

Soliciting Human Feedbacks to Improve REGEN. In many cases, there may exist difficult examples where the classifier and the retrieval model do not agree with each other. To enable the model to learn on these hard examples, *active learning* can be adopted to solicit human annotations (Yuan et al., 2020; Yu et al., 2022a,c) or instructions (Peng et al., 2023; Zhang et al., 2022b,a) to further improve the model performance.

Collaboration with Large Language Models. There are many other potential ways to incorporate black-box large language models into REGEN beyond our experiments. For instance, large language models can be used to *rerank* the top retrieved documents (Ma et al., 2023) or generate *augmented examples* for classifiers (Møller et al., 2023). On the other hand, REGEN can be integrated into the training set synthesis for language models when the labeled dataset is inaccessible (Zhang et al., 2023). It is still an open question on how to harness large language models for dataset generation in an efficient and effective way.

6.2 Conclusion

In this paper, we propose a framework REGEN for zero-shot text classification, which incorporates dense retrieval to synthesize task-specific training sets via retrieving class-relevant documents from the generic unlabeled corpus with verbalizers. We

further propose two simple while effective strategies to progressively improve the quality of the curated dataset. The effectiveness of REGEN is validated on nine benchmark datasets with an average gain of 4.3%. Further qualitative analysis justify the better quality of datasets generated by REGEN over baselines under multiple criteria.

Limitations

Our method REGEN is a general framework for zero-shot text classification. In this work, we aim to first bring in simple and intuitive way to justify the power of unsupervised dense retrieval for zero-shot learning. Effective as it is, there is still much room for improvements, including designing better objectives for pretraining R_θ as well as better strategies for removing noisy training data (Lang et al., 2022; Xu et al., 2023). How to improve these components is an important line of future work.

Besides, our experiment results are all based on BERT_{base} sized models. Although REGEN performs on par with or better than previous dataset generation methods using giant NLG models, it remains unknown to us how the benefit of REGEN scales with more parameters for both R_θ and C_ϕ .

Also, we point out that this work focuses on zero-shot *text classification* with task-specific verbalizers and unlabeled generic corpus, thus it can be nontrivial to adapt our framework to other tasks such as Natural Language Inference (NLI) as well as low-resource tasks where even the unlabeled generic corpus can be hard to collect. Extending REGEN to these settings will reduce the annotation burden under more challenging scenarios.

Ethics Statement

One potential risk of applying REGEN is that the generic corpus used in our experiments may contain harmful information as they were crawled from the Internet that are only filtered with some rules (Gehman et al., 2020). As a result, they may contain text exhibiting biases that are undesirable for target tasks. To alleviate this issue, we recommend the potential users to first use bias reduction and correction techniques (Schick et al., 2021) to remove biased text from the corpus to mitigate the risks of the curated dataset.

Acknowledgements

We thank anonymous reviewers for their feedback. This work was supported by NSF IIS-2008334, IIS-2106961, and CAREER IIS-2144338.

References

- Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millan, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego De Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, Tom Hennigan, Saffron Huang, Loren Maggioni, Chris Jones, Albin Cassirer, Andy Brock, Michela Paganini, Geoffrey Irving, Oriol Vinyals, Simon Osindero, Karen Simonyan, Jack Rae, Erich Elsen, and Laurent Sifre. 2022. Improving language models by retrieving from trillions of tokens. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 2206–2240. PMLR.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Xiang Chen, Lei Li, Ningyu Zhang, Xiaozhuan Liang, Shumin Deng, Chuanqi Tan, Fei Huang, Luo Si, and Huajun Chen. 2022. [Decoupling knowledge from memorization: Retrieval-augmented prompt learning](#). In *Advances in Neural Information Processing Systems*.
- Arman Cohan, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel Weld. 2020. [SPECTER: Document-level representation learning using citation-informed transformers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2270–2282, Online. Association for Computational Linguistics.
- Hejie Cui, Jiaying Lu, Yao Ge, and Carl Yang. 2022. How can graph neural networks help document retrieval: A case study on cord19 with concept map generation. In *Advances in Information Retrieval: 44th European Conference on IR Research, ECIR 2022, Stavanger, Norway, April 10–14, 2022, Proceedings, Part II*, pages 75–83. Springer.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- Luyu Gao and Jamie Callan. 2021. [Condenser: a pre-training architecture for dense retrieval](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 981–993, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Luyu Gao and Jamie Callan. 2022. [Unsupervised corpus aware language model pre-training for dense passage retrieval](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2843–2853, Dublin, Ireland. Association for Computational Linguistics.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021a. [Making pre-trained language models better few-shot learners](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, Online. Association for Computational Linguistics.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021b. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. 2020. [RealToxicityPrompts: Evaluating neural toxic degeneration in language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3356–3369, Online. Association for Computational Linguistics.
- Ariel Gera, Alon Halfon, Eyal Shnarch, Yotam Perlitz, Liat Ein-Dor, and Noam Slonim. 2022. [Zero-shot text classification with self-training](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1107–1119, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Daniel Gillick, Sayali Kulkarni, Larry Lansing, Alessandro Presta, Jason Baldrige, Eugene Ie, and Diego Garcia-Olano. 2019. [Learning dense representations for entity retrieval](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 528–537, Hong Kong, China. Association for Computational Linguistics.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Ruining He and Julian McAuley. 2016. [Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering](#). In *Proceedings of*

- the 25th International Conference on World Wide Web*, page 507–517.
- Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Jingang Wang, Juanzi Li, Wei Wu, and Maosong Sun. 2022. [Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer for text classification](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2225–2240, Dublin, Ireland. Association for Computational Linguistics.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022a. [Unsupervised dense information retrieval with contrastive learning](#).
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2022b. [Few-shot learning with retrieval augmented language models](#). *arXiv preprint arXiv:2208.03299*.
- Zhengbao Jiang, Luyu Gao, Jun Araki, Haibo Ding, Zhiruo Wang, Jamie Callan, and Graham Neubig. 2022. [Retrieval as attention: End-to-end learning of retrieval and reading within a single transformer](#). *arXiv preprint arXiv:2212.02027*.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2021. [Billion-scale similarity search with gpus](#). *IEEE Transactions on Big Data*, 7(3):535–547.
- Xuan Kan, Hejie Cui, and Carl Yang. 2021. [Zero-shot scene graph relation prediction through common-sense knowledge integration](#). In *Machine Learning and Knowledge Discovery in Databases. Research Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part II 21*, pages 466–482. Springer.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020. [Generalization through memorization: Nearest neighbor language models](#). In *International Conference on Learning Representations*.
- Hunter Lang, Aravindan Vijayaraghavan, and David Sontag. 2022. [Training subset selection for weak supervision](#). In *Advances in Neural Information Processing Systems*.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent retrieval for weakly supervised open domain question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy. Association for Computational Linguistics.
- Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N Mendes, Sebastian Hellmann, Mohamed Morsey, et al. 2015. [Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia](#). *Semantic web*, 6(2):167–195.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Chen Liang, Yue Yu, Haoming Jiang, Siawpeng Er, Ruijia Wang, Tuo Zhao, and Chao Zhang. 2020. [Bond: Bert-assisted open-domain named entity recognition with distant supervision](#). In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '20*, page 1054–1064, New York, NY, USA. Association for Computing Machinery.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A Robustly Optimized BERT Pretraining Approach](#). *arXiv preprint arXiv:1907.11692*.
- Xueguang Ma, Xinyu Zhang, Ronak Pradeep, and Jimmy Lin. 2023. [Zero-shot listwise document reranking with a large language model](#). *arXiv preprint arXiv:2305.02156*.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. [Learning word vectors for sentiment analysis](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150. Association for Computational Linguistics.
- 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, RecSys '13*, page 165–172, New York, NY, USA. Association for Computing Machinery.
- Yu Meng, Jiaxin Huang, Guangyuan Wang, Zihan Wang, Chao Zhang, Yu Zhang, and Jiawei Han. 2020a. [Discriminative topic mining via category-name guided text embedding](#). In *Proceedings of The Web Conference 2020, WWW '20*, page 2121–2132, New York, NY, USA. Association for Computing Machinery.
- Yu Meng, Jiaxin Huang, Yu Zhang, and Jiawei Han. 2022. [Generating training data with language models: Towards zero-shot language understanding](#). In *Advances in Neural Information Processing Systems*.
- Yu Meng, Yunyi Zhang, Jiaxin Huang, Chenyan Xiong, Heng Ji, Chao Zhang, and Jiawei Han. 2020b. [Text classification using label names only: A language model self-training approach](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural*

- Language Processing (EMNLP)*, pages 9006–9017, Online. Association for Computational Linguistics.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the role of demonstrations: What makes in-context learning work? *arXiv preprint arXiv:2202.12837*.
- Anders Giovanni Møller, Jacob Aarup Dalsgaard, Arianna Pera, and Luca Maria Aiello. 2023. Is a prompt and a few samples all you need? using gpt-4 for data augmentation in low-resource classification tasks. *arXiv preprint arXiv:2304.13861*.
- Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. 2019. When does label smoothing help? *Advances in neural information processing systems*, 32.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Bo Pang and Lillian Lee. 2005. [Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics*, pages 115–124, Ann Arbor, Michigan. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Baolin Peng, Michel Galley, Pengcheng He, Hao Cheng, Yujia Xie, Yu Hu, Qiuyuan Huang, Lars Liden, Zhou Yu, Weizhu Chen, and Jianfeng Gao. 2023. [Check your facts and try again: Improving large language models with external knowledge and automated feedback](#).
- Ethan Perez, Douwe Kiela, and Kyunghyun Cho. 2021. [True few-shot learning with language models](#). In *Advances in Neural Information Processing Systems*.
- Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, Vassilis Plachouras, Tim Rocktäschel, and Sebastian Riedel. 2021. [KILT: a benchmark for knowledge intensive language tasks](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2523–2544, Online. Association for Computational Linguistics.
- Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaid Harchaoui. 2021. [MAUVE: Measuring the gap between neural text and human text using divergence frontiers](#). In *Advances in Neural Information Processing Systems*.
- Stephen Robertson, Hugo Zaragoza, and Michael Taylor. 2004. [Simple bm25 extension to multiple weighted fields](#). In *Proceedings of the Thirteenth ACM International Conference on Information and Knowledge Management*, page 42–49, New York, NY, USA. Association for Computing Machinery.
- Ohad Rubin, Jonathan Herzig, and Jonathan Berant. 2022. [Learning to retrieve prompts for in-context learning](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2655–2671, Seattle, United States. Association for Computational Linguistics.
- Devendra Singh Sachan, Siva Reddy, William L. Hamilton, Chris Dyer, and Dani Yogatama. 2021. [End-to-end training of multi-document reader and retriever for open-domain question answering](#). In *Advances in Neural Information Processing Systems*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, et al. 2022. [Multitask prompted training enables zero-shot task generalization](#). In *International Conference on Learning Representations*.
- Timo Schick and Hinrich Schütze. 2021a. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021b. [Generating datasets with pretrained language models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6943–6951, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Timo Schick, Sahana Udupa, and Hinrich Schütze. 2021. Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in nlp. *Transactions of the Association for Computational Linguistics*, 9:1408–1424.
- Jiaming Shen, Wenda Qiu, Yu Meng, Jingbo Shang, Xiang Ren, and Jiawei Han. 2021. [TaxoClass: Hierarchical multi-label text classification using only class names](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4239–4249, Online. Association for Computational Linguistics.
- Weijia Shi, Julian Michael, Suchin Gururangan, and Luke Zettlemoyer. 2022. Nearest neighbor zero-shot inference. *arXiv preprint arXiv:2205.13792*.

- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Si Sun, Chenyan Xiong, Yue Yu, Arnold Overwijk, Zhiyuan Liu, and Jie Bao. 2022a. [Reduce catastrophic forgetting of dense retrieval training with teleportation negatives](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6639–6654, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yi Sun, Yu Zheng, Chao Hao, and Hangping Qiu. 2022b. [NSP-BERT: A prompt-based few-shot learner through an original pre-training task — next sentence prediction](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3233–3250, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Derek Tam, Rakesh R. Menon, Mohit Bansal, Shashank Srivastava, and Colin Raffel. 2021. [Improving and simplifying pattern exploiting training](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4980–4991, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mozes van de Kar, Mengzhou Xia, Danqi Chen, and Mikel Artetxe. 2022. [Don’t prompt, search! mining-based zero-shot learning with language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7508–7520, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hanan Hajishirzi. 2022. [Self-instruct: Aligning language model with self generated instructions](#). *arXiv preprint arXiv:2212.10560*.
- Zihan Wang, Dheeraj Mekala, and Jingbo Shang. 2021. [X-class: Text classification with extremely weak supervision](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3043–3053, Online. Association for Computational Linguistics.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022. [Finetuned language models are zero-shot learners](#). In *International Conference on Learning Representations*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, et al. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *arXiv preprint arXiv:1910.03771*.
- Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul N. Bennett, Junaid Ahmed, and Arnold Overwijk. 2021. [Approximate nearest neighbor negative contrastive learning for dense text retrieval](#). In *International Conference on Learning Representations*.
- Jinxi Xu and W. Bruce Croft. 2017. [Query expansion using local and global document analysis](#). *SIGIR Forum*, 51(2):168–175.
- Ran Xu, Yue Yu, Hejie Cui, Xuan Kan, Yanqiao Zhu, Joyce C. Ho, Chao Zhang, and Carl Yang. 2023. [Neighborhood-regularized self-training for learning with few labels](#). In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence*.
- Yiben Yang, Chaitanya Malaviya, Jared Fernandez, Swabha Swayamdipta, Ronan Le Bras, Ji-Ping Wang, Chandra Bhagavatula, Yejin Choi, and Doug Downey. 2020. [Generative data augmentation for common-sense reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1008–1025, Online. Association for Computational Linguistics.
- Jiacheng Ye, Jiahui Gao, Jiangtao Feng, Zhiyong Wu, Tao Yu, and Lingpeng Kong. 2022a. [Progen: Progressive zero-shot dataset generation via in-context feedback](#). *arXiv preprint arXiv:2210.12329*.
- Jiacheng Ye, Jiahui Gao, Qintong Li, Hang Xu, Jiangtao Feng, Zhiyong Wu, Tao Yu, and Lingpeng Kong. 2022b. [ZeroGen: Efficient zero-shot learning via dataset generation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11653–11669, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Wenpeng Yin, Jamaal Hay, and Dan Roth. 2019. [Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3914–3923, Hong Kong, China. Association for Computational Linguistics.
- HongChien Yu, Chenyan Xiong, and Jamie Callan. 2021. [Improving query representations for dense retrieval with pseudo relevance feedback](#). In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management, CIKM ’21*, page 3592–3596, New York, NY, USA. Association for Computing Machinery.
- Yue Yu, Ling kai Kong, Jieyu Zhang, Rongzhi Zhang, and Chao Zhang. 2022a. [AcTune: Uncertainty-based active self-training for active fine-tuning of pretrained](#)

- language models. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1422–1436, Seattle, United States. Association for Computational Linguistics.
- Yue Yu, Chenyan Xiong, Si Sun, Chao Zhang, and Arnold Overwijk. 2022b. [COCO-DR: Combating the distribution shift in zero-shot dense retrieval with contrastive and distributionally robust learning](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1462–1479, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yue Yu, Rongzhi Zhang, Ran Xu, Jieyu Zhang, Jiaming Shen, and Chao Zhang. 2022c. Cold-start data selection for few-shot language model fine-tuning: A prompt-based uncertainty propagation approach. *arXiv preprint arXiv:2209.06995*.
- Michelle Yuan, Hsuan-Tien Lin, and Jordan Boyd-Graber. 2020. [Cold-start active learning through self-supervised language modeling](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7935–7948, Online. Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. Defending against neural fake news. *Advances in neural information processing systems*, 32.
- Jieyu Zhang, Yue Yu, Yinghao Li, Yujing Wang, Yaming Yang, Mao Yang, and Alexander Ratner. 2021. [WRENCH: A comprehensive benchmark for weak supervision](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Rongzhi Zhang, Jiaming Shen, Tianqi Liu, Jialu Liu, Michael Bendersky, Marc Najork, and Chao Zhang. 2023. [Do not blindly imitate the teacher: Using perturbed loss for knowledge distillation](#).
- Rongzhi Zhang, Rebecca West, Xiquan Cui, and Chao Zhang. 2022a. Adaptive multi-view rule discovery for weakly-supervised compatible products prediction. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4521–4529.
- Rongzhi Zhang, Yue Yu, Pranav Shetty, Le Song, and Chao Zhang. 2022b. [Prompt-based rule discovery and boosting for interactive weakly-supervised learning](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 745–758, Dublin, Ireland. Association for Computational Linguistics.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28:649–657.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*, pages 12697–12706. PMLR.
- Ruiqi Zhong, Kristy Lee, Zheng Zhang, and Dan Klein. 2021. [Adapting language models for zero-shot learning by meta-tuning on dataset and prompt collections](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2856–2878, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. Texus: A benchmarking platform for text generation models. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1097–1100.
- Yuchen Zhuang, Yinghao Li, Junyang Zhang, Yue Yu, Yingjun Mou, Xiang Chen, Le Song, and Chao Zhang. 2022. [ReSel: N-ary relation extraction from scientific text and tables by learning to retrieve and select](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 730–744, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

A Verbalizers and Templates for Datasets

The verbalizers and templates of datasets are shown in table 9.

B Corpus

We select three types of corpus, i.e. WIKI (Petroni et al., 2021), subsets of REVIEWS (He and McAuley, 2016) and REALNEWS (Zellers et al., 2019) to form the corpus \mathcal{D} . We manually remove documents less than 10 words as we observe that these documents do not contain informative content. The detailed information is shown in table 10.

C Baselines

We consider multiple baselines for zero-shot text classification. The details of these baselines are described as follows. We use * to denote baselines with extra resources or large language models.

Zero-shot Inference Methods These methods directly inference over the test set for prediction.

- **NSP-BERT** (Sun et al., 2022b): It uses the next sentence prediction (NSP) task to perform zero-shot learning. Specifically, it construct prompts for each labels, and use the PLM with the NSP head as the indicator.
- **Prompt** (Schick and Schütze, 2021a): It uses the original masked language modeling (MLM) objective with category-specific verbalizers to infer the true label of each sentence.
- **KNN-Prompt** (Shi et al., 2022): It improves zero-shot prompting by retrieving relevant information from an additional heterogeneous corpus, which achieves better coverage of the verbalizers.
- **KPT*** (Hu et al., 2022): It uses additional knowledge bases (e.g. WordNet) to expand the label word space for verbalizers, for improving prompt-based learning.
- **GPT-3*** (Brown et al., 2020): It adopts GPT-3 for zero-shot learning. We use the contextual calibration (Zhao et al., 2021) by default as it can improve the zero-shot prediction accuracy.

Transfer-learning based Inference Methods

- **TE-NLI*** (Yin et al., 2019): It uses the model fine-tuned on NLI tasks to perform zero-shot classification.

- **NLI-ST*** (Gera et al., 2022): It uses self-training to finetune the model on additional unlabeled task-specific corpus.

We are aware that there exist some other models for generic zero-shot learning on NLP such as FLAN (Wei et al., 2022) and T0 (Sanh et al., 2022), we do not compare with them since they leverage the labeled data from some of the datasets evaluated in this work (e.g. AGNews, IMDB, according to their original paper). It is thus inappropriate to use them under the true zero-shot learning setting, since such models can have unfair advantages due to access to related data during pre-training.

Weakly-supervised Learning Methods This line of methods is close to the general zero-shot learning in the sense that it does not rely on any labeled examples for classification (Shen et al., 2021; Liang et al., 2020; Zhang et al., 2021). Instead, it leverages class-specific verbalizers as well as *task-specific* unlabeled data as *weak supervision* for classification.

- **LOTClass*** (Meng et al., 2020b): It first matches the label name with the corpus to find category-indicative words, then trains the model to predict their implied categories with self-training.
- **X-Class*** (Wang et al., 2021): It estimates class representations by adding the most similar word to each class, then obtains the document representation with weighted average of word representations. Finally, the most confidence words are selected to fine-tune the classifier.

Note that we present the results for the two methods, but mainly for *reference* purposes as the setting between these approaches and our work is different.

Dataset Generation Methods These methods generates specific datasets for zero-shot learning. Note that we use the same pretrained RoBERTa-base model as the classifier and use the same label smoothing loss for fine-tuning.

- **SuperGen** (Meng et al., 2022): It is one of the representative methods for using large natural language generation models (NLG) for zero-shot learning. It first uses the NLG model to generate training data with prompts, then selects data with highest generation probability for fine-tuning.

Task	Verbalizers	Template used for Retrieval	Template used for Prompting
AG News	politics, sports business, technology	[VERB] News.	The category of x^b is [VERB].
DBPedia	company, school, artist, athlete politics, transportation, building, river/mountain/lake, village, animal, plant, album, film, book	[VERB]	$x^a x^b?$ The category of x^a is [VERB].
Yahoo	society, science, health, school computer, sports, business, music, family, politics	[VERB]	$x^a x^b?$ The category of x^a is [VERB].
NYT	business, politics, sports health, education, estate art, science, technology	[VERB] News.	The category of x^b is [VERB].
Sentiment	great bad	It was a [VERB] movie.	It was a [VERB] movie.. x^b .

Table 9: The format of verbalizers and the template used for retrieval and prompting. We use the prompt formats provided in prior works (Schick and Schütze, 2021a; Hu et al., 2022). The [VERB] stands for the verbalizers. x^a stands for the title (only exist in DBPedia and Yahoo) and x^b stands for the body of the target document.

Corpus	Size	Size after Pre-processing
Wiki (Petroni et al., 2021)	6M	6M
News (Zellers et al., 2019)	11.9M	6M
Reviews (He and McAuley, 2016)	24.0M	4M

Table 10: The information about the general corpus \mathcal{D} used in this study.

- **Mining** (van de Kar et al., 2022): It uses regular expressions with category-related keywords to mine samples (the *next* sentences of the matched text) from the corpus to generate training data. Then, it uses the zero-shot prompting to filter the noisy sample and fine-tune another classification model on the filtered dataset. For fair comparison, we use the same corpus \mathcal{D} , prompt format as ours for zero-shot learning, note that these often result in better performance.

The comparison of REGEN with other methods within this category (e.g. (Ye et al., 2022a,b)) is shown in Appendix F.

D Implementation Details

D.1 Implementation Details for Baselines

For *zero-shot inference* methods, we directly use the numbers from the original papers if available, and reimplement Dataless and Prompt on our own. From our experiments, we observe that the numbers reported in van de Kar et al. (2022) is much lower than our reimplemented prompt-based zero-shot learning results, for reasons unknown to us.

For *transfer-learning based zero-shot inference* methods, we use the same verbalizer as REGEN

and the prompt template provided from the authors for inference with the released pretrained models.

For *weakly-supervised learning* and *zero-shot dataset generation* methods, we use the code released by the authors with the optimal hyperparameters reported in the corresponding paper if available. As the code for (van de Kar et al., 2022) is **not publicly available**, we reimplement this method based on the information from the paper. If fine-tuning is involved, we use the same pretrained RoBERTa-base as the classifier C_ϕ with the label smoothing strategy for fair comparison.

D.2 Implementation Details for REGEN

Table 11 lists the hyperparameters used for REGEN. Note that we keep them *same* across all tasks without any further tuning. Under the zero-shot learning setting, there is *no validation set* available. For each task, we follow (Ye et al., 2022b) to use a portion (e.g., 10%) of the pseudo dataset as the validation set for model selection. If the total number of the training data for a specific category exceeds 3000, we randomly sample a subset with 3000 samples for that category.

lr_{ft}	lr_{cl}	bsz_{ft}	bsz_{cl}	$ \tilde{\mathcal{T}}^T $	E_1	E_2	T	α	τ	(k_1, k_2, k_3)
1e-5	1e-4	32	400	3,000	5	5	3	0.1	1	(100, 20, 20) for sentiment and (50, 10, 10) for topics

Table 11: Hyperparameters on different tasks (they are kept same for all tasks). lr_{ft} : Learning rate for fine-tuning; lr_{cl} : Learning rate for unsupervised contrastive learning; bsz_{ft} : Batch size; bsz_{cl} : Batch size for unsupervised contrastive learning; $|\tilde{\mathcal{T}}^T|$: Maximum number of selected training data per class after the final retrieval round; E_1 : Number of epochs for fine-tuning; E_2 : Number of epochs for contrastive learning; T : Number of retrieval rounds, α : Parameter for label smoothing ; τ : Temperature parameter for contrastive learning; (k_1, k_2, k_3) : Parameter k used in ANN in each round.

Task	Template ID	Verbalizers
AGNews	#0 (Original)	politics, sports, business, technology
	1	world, football, stock, science
	2	international, basketball, financial, research
	3	global, tennis, profit, chemical
Sentiment	#0 (Original)	great, bad
	1	good, awful
	2	awesome, terrible
	3	incredible, horrible

Table 12: Different verbalizers used for experiments in section 5.7.

D.3 Number of Parameters in REGEN

The retrieval model R_θ uses BERT-base-uncased as the backbone with 110M parameters, and the classification model C_ϕ uses RoBERTa-base as the backbone with 125M parameters.

D.4 Computation Environment

All experiments are conducted on *CPU*: Intel(R) Core(TM) i7-5930K CPU @ 3.50GHz and *GPU*: NVIDIA GeForce RTX A5000 GPUs using python 3.8 and Pytorch 1.10.

E Additional Information on Experiments Setups

E.1 Setup for Fine-tuning C_ϕ with Few Labeled Examples

Under the few-shot setting, we follow (Meng et al., 2022) to split the data into two parts: half the data as training set, and the remaining as the validation set. When a few labeled samples are available, we first fine-tune the classifier C_ϕ on the few-shot training set (denoted as C_ϕ^{init}), and use C_ϕ^{init} to remove the noisy instances with the method in Eq. 7 for both our method and baselines. Then, we continue fine-tuning the classifier on the generated data.

E.2 Setup for Zero-shot Learning with Different Verbalizers

We list the set of verbalizers used for section 5.7 in table 12.

E.3 Setup for Large Language Models for Verbalizer Expansion

For verbalizer expansion, we use GPT-4 (OpenAI, 2023) as the LLM backbone, and the prompt format is shown in the followings:

Suppose you are asked to perform text classification with the following labels. Can you generate 10 relevant keywords for each of the categories?

By inputting the verbalizers of each class into the chatbox, the LLM can output a series of keywords to enrich the verbalizer. After obtaining the keywords, we manually remove keywords that occur in more than one category, and the remaining keywords will be used for retrieval.

F Comparison with Recent Baselines

We provide additional empirical studies to compare REGEN with some recent works. As (Ye et al., 2022a,b) use a smaller PLM, namely DistillBERT (Sanh et al., 2019) for their experiments, we use the same DistillBERT encoder to finetune our model and several baselines (e.g. Mining (van de Kar et al., 2022) and SuperGen (Meng et al., 2022)). The result is shown in table 13.

Overall, we observe that REGEN outperforms most of these baselines with DistillBERT as the classifier. It achieves competitive performance with ProGen, which relies on several additional techniques including influence estimation, multi-round in-context feedback from a billion-scale language model, and noise-robust loss functions. Note that these techniques are orthogonal to our method, and can be potentially integrated with REGEN for better performance.

Method/Dataset	IMDB	SST-2	Rotten Tomato	Elec	Yelp	Avg.
Prompting*	77.31	82.63	78.66	78.03	80.30	79.39
ZeroGen* (Ye et al., 2022b)	80.41	82.77	78.36	85.35	87.84	82.94
ProGen* (Ye et al., 2022a)	84.12	87.20	82.86	89.00	89.39	86.51
SuperGen (Meng et al., 2022)	84.58	86.70	79.08	90.58	89.98	86.18
Mining (van de Kar et al., 2022)	77.36	80.73	76.73	85.87	90.36	82.21
REGEN	87.84	85.32	81.42	89.83	89.00	86.68

Table 13: Results with recent baselines using DistillBERT (Sanh et al., 2019) as C_ϕ . *: Results are copied from the previous papers (Ye et al., 2022a,b).

Task	Datasets	Performance of REGEN	Fully-supervised Performance	Δ Performance Gap	Lexical Similarity
Topic	AG News	85.0	94.6	10.0%	0.427
	DBPedia	87.6	99.2	11.2%	0.566
	Yahoo	59.4	76.8	17.4%	0.362
	NYT	74.5	88.2	15.5%	0.530
Sentiment	IMDB	89.9	94.4	4.5%	0.497
	MR	82.5	91.3	8.8%	0.306
	SST-2	88.9	96.2	7.3%	0.296
	Amazon	92.3	95.4	3.1%	0.714
	Yelp	93.0	97.2	4.2%	0.408

Table 14: The detailed value for the performance gap and the lexical similarity between the task-specific corpus and the general-domain corpus \mathcal{D} .

G More Details on Performance Gaps and Lexical Similarities

G.1 Calculating the Similarity between the Corpus and Target Tasks

We use the weighted Jaccard similarity $J(T, \mathcal{D})$ to measure distribution similarities between the corpus \mathcal{D} and the target task T , described as follows: Denote C_k as the frequency of word k in the corpus \mathcal{D} and T_k for the target task T respectively. The weighted Jaccard similarity $J(T, \mathcal{D})$ is defined as:

$$J(T, \mathcal{D}) = \frac{\sum_k \min(C_k, T_k)}{\sum_k \max(C_k, T_k)}, \quad (8)$$

where the sum is over all unique words k present in \mathcal{D} and T .

G.2 The Performance Gap and Lexical Similarity for All Datasets

The details for the performance gap as well as the lexical similarity to the general-domain corpus are shown in Table 14.

H Additional Per-task Results

We show the results for each task in this section. Specifically, we present the performance of REGEN and its variation of without the filtering step in Fig. 6; we present the performance of REGEN with different dense retrieval models as R_θ in Fig. 7; we illustrate the performance under different volume

Dataset	Verbalizer Group	Mining	SuperGen	REGEN
Yelp	# 0 (Original)	92.3	93.6	93.0
	# 2	85.4	91.6	91.9
	# 3	93.4	91.2	94.5
	# 4	93.2	93.2	92.8
	Avg. \pm Std.	91.1 \pm 3.8	92.4 \pm 1.2	93.1\pm1.1
Amazon	# 0 (Original)	92.0	91.0	92.3
	# 1	86.8	90.6	91.0
	# 2	91.4	88.9	93.1
	# 3	90.7	91.5	92.0
	Avg. \pm Std.	90.2 \pm 2.3	90.5 \pm 1.1	92.1\pm0.8
MR	# 0 (Original)	79.7	81.9	82.5
	# 1	79.5	80.8	83.6
	# 2	82.3	79.1	85.2
	# 3	81.6	82.2	83.1
	Avg. \pm Std.	80.8 \pm 1.3	81.0 \pm 1.4	83.6\pm1.2
SST-2	# 0 (Original)	85	88.6	88.9
	# 1	84.2	86.6	88.2
	# 2	87.8	85.4	89.5
	# 3	86.7	86.8	88.4
	Avg. \pm Std.	85.9 \pm 1.6	86.8 \pm 1.3	88.8\pm0.6

Table 15: Results with different verbalizers on other sentiment analysis datasets.

of training data for REGEN and baselines in Fig. 8; we demonstrate the effect of different corpus \mathcal{D} on the final performance in Fig. 9. Besides, in table 15 we illustrate the performance of REGEN and baselines on all sentiment analysis datasets; in table 16, the automatic evaluation results for all datasets are shown.

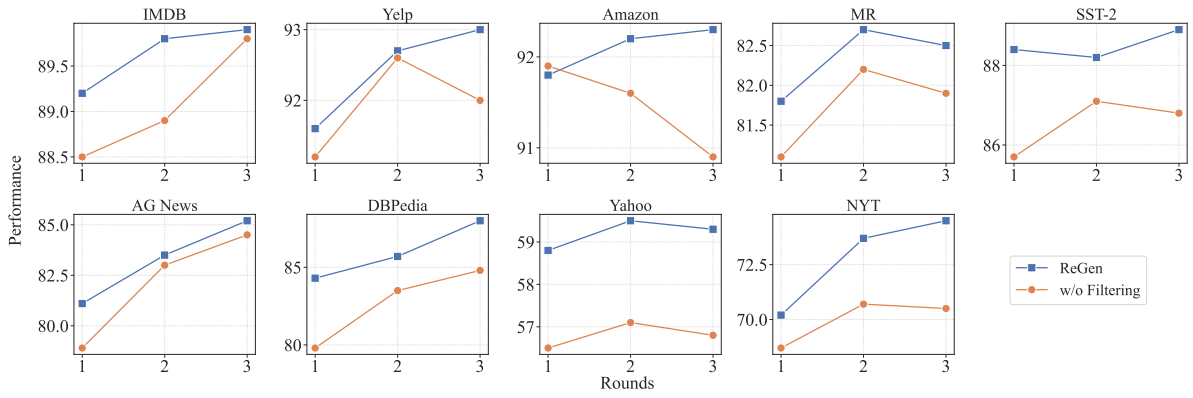


Figure 6: Effect of filtering, per task results.




Figure 7: Comparisons of different dense retrieval models, per task results.

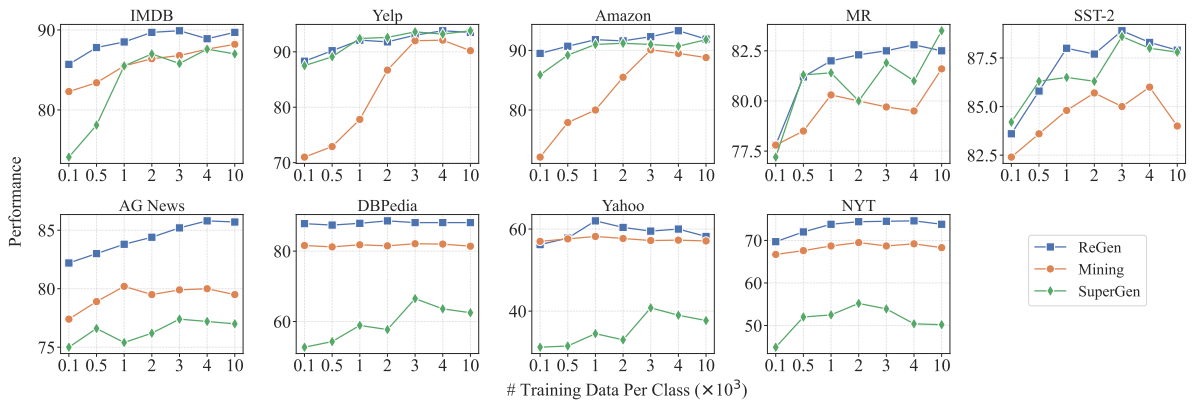


Figure 8: Performance of the different amount of training data, per task results.

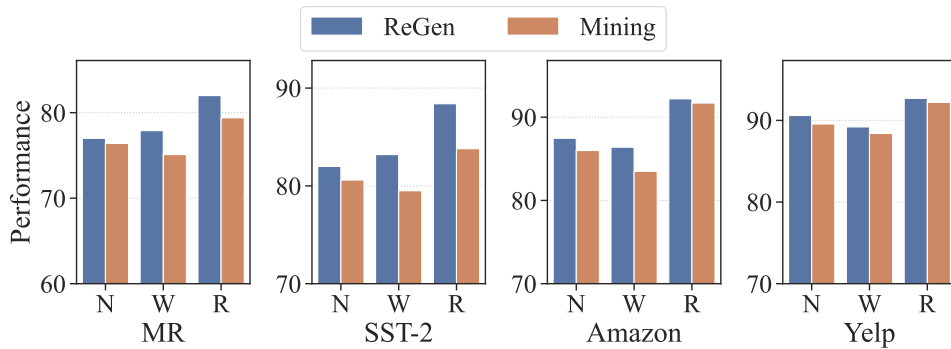


Figure 9: Performance of REGEN using different subsets of corpus on other sentiment classification tasks.

Dataset	Metrics	Mining	SuperGen	REGEN
Sentiment	Correctness (\uparrow)	0.815	0.971	0.986
	Diversity (\downarrow)	0.144	0.915	0.361
	Distribution Sim. (\uparrow)	0.856	0.803	0.865
AG News	Correctness (\uparrow)	0.746	0.649	0.805
	Diversity (\downarrow)	0.117	0.818	0.330
	Distribution Sim. (\uparrow)	0.799	0.687	0.686
DBPedia	Correctness (\uparrow)	0.791	0.516	0.909
	Diversity (\downarrow)	0.223	0.765	0.377
	Distribution Sim. (\uparrow)	0.874	0.662	0.920
NYT	Correctness (\uparrow)	0.730	0.811	0.893
	Diversity (\downarrow)	0.100	0.717	0.342
	Distribution Sim. (\uparrow)	0.511	0.643	0.622
Yahoo	Correctness (\uparrow)	0.771	0.518	0.832
	Diversity (\downarrow)	0.089	0.768	0.335
	Distribution Sim. (\uparrow)	0.810	0.602	0.797

Table 16: Automatic evaluation results on all datasets. Note that we only generate one dataset for all sentiment analysis tasks.

I Details for Quality Analysis

I.1 Automatic Evaluation

We provide the details for automatic measurements of the dataset quality as follows.

For *correctness*, we first fine-tune a RoBERTa-Large model on the original dataset⁸, and use the fine-tuned model as an oracle to evaluate the correctness of the synthetic dataset.

For *diversity*, we use the self-BLEU (Zhu et al., 2018), which computes the BLEU-4 score of each generated text with other generations in the dataset as references, as the metric. Note that for self-BLEU, a *lower* score implies higher diversity.

Besides, we use MAUVE (Pillutla et al., 2021) with the default hyperparameter settings to measure the *distribution similarity*. MAUVE is originally proposed for comparing the learnt distribution of a text generation model and the distribution of human-written text, and we adapt MAUVE to measure the similarity between the distribution of the synthetic dataset and the real dataset. A higher value indicates that the distribution of the synthetic dataset and the real dataset is closer, thus the quality of the synthetic dataset is higher.

I.2 Human Evaluation

Apart from the automatic evaluation, we also perform human evaluation to manually evaluate the quality of the synthetic dataset. We ask four volunteer students from our institute (approved by the ethics review board) for participation. For human evaluation, the evaluation form is listed as below.

⁸For sentiment analysis, we combine the training set of five datasets together as the final training set.

- **Correctness:** Whether the text is relevant to the corresponding label?
 - 2: Accurate: The content is accurate for the label.
 - 1: Related: The content is related but not accurate for the label.
 - 0: Not relevant: The content is not relevant to the label.
- **Informativeness:** Whether the text is fluent and similar to human-generated text?
 - 2: Very Informative: The text is very informative and similar to human generated text.
 - 1: Partially Informative: The text is partially informative and somewhat close to human generated text.
 - 0: Not Informative: The text is not fluent/informative at all.
- **Diversity:** Whether the generated text within each class is diverse enough?
 - 2: Diverse: The documents cover different topics related to the label.
 - 1: Partially Diverse: There exists some overlapping documents.
 - 0: Not Diverse: All the documents look similar.

J Case Studies

We present some examples of the curated dataset in the table 17 and 18. Note that *filtered* means the data is first retrieved by R_θ but is later identified as incorrect sample by the classifier. Overall, we observe that the dataset of SuperGen contains similar sentences across samples (*e.g.* a great example of the kind of movie for sentiment analysis datasets), and the mining-based approach often contains documents that are less informative (*e.g.* World famous hapuna beach is just minutes away for topic classification). In contrast, REGEN goes beyond the exact keyword matching and can retrieve diverse and informative documents. Moreover, the noisy samples can be filtered through self-consistency regularization. These cases corroborate the result on both automatic and human evaluation, and justify the higher quality of the dataset curated by us.

We also demonstrate the retrieved examples over different rounds in table 19 and 20. Note that examples shown in the 2nd and 3rd round are retrieved

Method	Label	Generated Text
SuperGen	positive	The film is a great example of the kind of movie that you can watch over and over.
	positive	The movie was very good and it had a lot of action in it. I would recommend this to anyone who likes action..
	positive	This film is a great example of the kind of movie that you can watch with your kids and not have to worry about anything inappropriate.
	negative	The film was a total waste of time. I would not recommend this movie to anyone.
	negative	This film is a waste of time. It has no plot and the acting was terrible. I would not recommend this movie to anyone.
	negative	This movie is not worth the time or money to watch it. It was a waste of my time.
Mining	positive	The dvd arrived very quick..
	positive	I can't stop playing them right now.
	positive	I cannot use it on the hardwood floors because i am afraid water might get down under them (they are very old but have been refinished.).
	negative	The plastic handle is not able to be taken apart so i don't know where the leak was exactly coming from.
	negative	Don't know this for sure, but it seems likely.
	negative	OK, this cd makes me sad.
REGEN	positive	Great I bought this toy for my son's 3rd birthday and only after 2 months he now sings the alphabet song all the time. It is a great education toy and also very durable.
	positive	After seeing the movie "12 Years A Slave," I wanted to read the book. The experience of watching the movie drew me into the story of Solomon Northup's life.
	positive	This is a must see film for all ages I would have given this film 10 stars if they would have let me. This is one of those films that somehow got overlooked in the theaters..
	positive (filtered)	Excellent but still not Perfect. Don't take my title or rating the wrong way. My experience with the first 2 Harry Potter Movies have been excellent, but in the 2nd movie, the Chamber of Secrets, A lot of parts were taken out...
	negative	Worst movie ever A good example of what is wrong with Hollywood today. I have never looked at my watch more times during a movie.
	negative	Bad book I did not like it. It is a bad story. Wolfs are not bad like in the story. Peter doesnt listen to his grandpa, so it is a bad example.
	negative	Silicon Valley... I do not like this game. The directions are hard to follow and I did not like the graphics at all.
	negative (filtered)	how can people dislike this charming film, this is very wonderful film that works for both audlts and kids.

Table 17: Example retrieved texts of REGEN and two baselines on synthetic dataset for sentiment analysis.

Method	Label	Generated Text
SuperGen	politics	The opinions expressed in this commentary are solely those of John Avlon.
	politics	TL;DR Correction of Sept 30 article on Pres Bush's visit to New York City, which misstated his role in campaign finance reform legislation that was signed into law by Gov George Pataki.
	sports	TL;DR Correction of Nov 12 article on New York Yankees pitcher Roger Clemens regarding his use of steroids; he is not using steroids and has never used them.
	sports	TL;DR Correction of Aug 25 article on New York Yankees player Mariano Rivera regarding his role in team's World Series victory over Arizona Diamondbacks.
	business	The company said it had hired the law firm of Paul, Weiss, Rifkind, Wharton & Garrison to conduct an independent investigation.
Mining	business	The company said it had hired the law firm of Debevoise & Plimpton to conduct an independent investigation.
	technology	TL;DR The National Science Foundation awarded \$32 million to the University of California, Berkeley, for research on how people use computers in their lives.
	technology	TL;DR The New York Times Magazine publishes its annual list of the 100 most influential people in science, technology, engineering or math.
	politics	World famous hapuna beach is just minutes away.
	politics	At the same time, we should not let our good fortune make us callous to the effect of suffering on most of the world population.
REGEN	sports	According to multiple sportsbooks, curry isn't even in the top-five likeliest mvp candidates for 2016-17.
	sports	Sky sports reported tonight chelsea have held talks over the former napoli manager's future.
	business	I am not stary-eyed about the news business 2014 and it is a business.
	business	Fostering a sense of autonomy amongst employees should be a central goal for all business leaders.
	technology	Notebook casing supplier catcher technology was forced to close one facility over environmental concerns, while iphone supplier pegatron was fined for spewing harmful gases during the manufacture of products.
REGEN	technology	Panaji: goa police in association with a bengaluru-based start-up has come up with a technology which can detect unauthorized drones.
	politics	The United Nations Human Rights Commissioner Navi Pillay has called for an international probe into war crimes committed in Sri Lanka during the final stages of its ethnic conflict, according to a media report on Sunday.
	politics	Police in Bolivia have rebelled against the government, abandoning their posts and marching through the streets along with protesters. It's a sign of growing anger over alleged voter fraud in last month's election. Protests since the poll have resulted in three deaths.
	politics (filtered)	An Australian in ASEAN. It sounds like the title of an innocent-abroad movie: the hero has adventures, blunders and embarrasses. But in the end Aussie charm and grit prevail; romance blossoms and the outsider becomes an insider..
	sports	Tom Brady and Bill Belichick likely will go down as the greatest quarterback/coach combo in NFL history, especially after winning their fifth Super Bowl together with a thrilling 34-28 overtime victory against the Atlanta Falcons in Super Bowl LI on Sunday night.
	sports	Manchester City's quest for four trophies continued with a 5-0 thrashing of Burnley to march into the FA Cup fifth round as League One Shrewsbury narrowly missed out on shocking Wolves in a 2-2 draw on Saturday.
	sports (filtered)	The growing scandal involving the new designer steroid THG gives sports fans one more thing other than sports to worry over. To be a sports fan is to get a constant education in subjects that don't necessarily interest you.
	business	THE HAGUE, Netherlands, March 14, 2019 /PRNewswire/ – Royal Dutch Shell plc RDS.A, +0.35% RDS.B, +0.19% filed its Annual Report on Form 20-F for the year ended December 31, 2018, with the U.S. Securities and Exchange Commission.
	business	Dimensions International Inc. has acquired Sentel Corp., creating a company that will have more than \$100 million in annual revenue. Terms of the deal were not disclosed.
	business (filtered)	Mercosur full members (Argentina, Brazil, Paraguay and Uruguay) rank poorly in the Forbes magazine annual Best Countries for Business, with the best listed, Chile and Peru, in positions 24 and 42, out of 134 countries surveyed worldwide.
REGEN	technology	SpaceX's next-generation rocket, the Starship, is 50 meters long and powered by three Raptor engines, creating a whopping 12,000 kN of thrust. It is designed to haul large amounts of cargo and eventually passengers into space, for missions to the moon and potentially to Mars and beyond as well.
	technology	Physicians that use the clinical reference tool, DynaMedTM from EBSCO Health, can now access the valuable, evidence-based content anywhere with the new DynaMed mobile app. The new app has been redesigned to make it easier and faster for physicians to find answers to clinical questions.
	technology	Cookson is science editor at the FT. He joined the newspaper in 1988 as technology editor and has also written about the chemical and pharmaceutical industries.
	technology (filtered)	Previously, he was the science and medical correspondent for BBC Radio.

Table 18: Example retrieved texts of REGEN and two baselines on the synthetic dataset for AG News.

directly using the concatenation of class-specific verbalizers and document from the previous round. The results indicate that REGEN can iteratively retrieve text that are semantically close to the documents from previous rounds.

Round	Label	Generated Text
1	positive	"Deceptions" was one of the best films I have seen in a long time. Stefanie Powers was excellent as Sabrina and Samantha. The rest of the cast was also very good.
	negative	I honestly have no idea what to say about this movie. It literally left me speechless... in a very, very not-good way.
2	positive	I saw the film last weekend and enjoyed it. From the point of view of movie craftsmanship, it's hard to go wrong with the talent combination of Steven Spielberg, Meryl Streep, Tom Hanks, and John Williams.
	negative	To be frank, it is a really bad movie. The cheap symbolism would make a junior high English teacher blush (including the title), and the lopsided view of racism in America was painfully and repeatedly portrayed.
3	positive	"Letting Go," with Sharon Gless and John Ritter, was a warm, funny and dramatic movie. I loved it. It was a fresh and wonderful romance.
	negative	First of all, I would like to say that I think the movie did an excellent job of following the events in the book. But they did a pretty bad job of leaving some crucial parts out of the movie. In the book, you get a pretty strong sense of the bond and relationship between the characters. In the movie, you don't really see that bond at all.

Table 19: Example retrieved texts of REGEN over three rounds for sentiment datasets.

Round	Label	Generated Text
1	politics	The UN voiced hope Monday that a meeting this week of a committee tasked with amending Syria's constitution can open the door to a broader political process for the war-ravaged country.
	sports	LaLiga may boast football superpowers Real Madrid and Barcelona but the league is keen to help other Spanish sports succeed too.
	business	Corporate America is slowly starting to give cash back to investors with dividends and buybacks. Companies are also spending cash on mergers.
	technology	Google said on Wednesday it had achieved a breakthrough in research, by solving a complex problem in minutes with a so-called quantum computer that would take today's most powerful supercomputer thousands of years to crack.
2	politics	The death toll in Eastern Ghouta stands at nearly 500, and it remains unclear how the sustained bombing campaign in the region will stop—despite a UN vote.
	sports	Barcelona continued their quest to win La Liga with a comfortable 3-0 victory over Leganes yesterday. Luis Suarez ended his goal drought with a brilliant brace before summer signing Paulinho got on the scoresheet late on.
	business	For many American companies today it is almost as if the recession never happened as executive incomes rise above pre-recession levels. According to Standard & Poor's 500 the average income of an executive in 2010 was \$9 million. That is 24 percent higher than it was the year prior.
	technology	Scientists claimed Wednesday to have achieved a near-mythical state of computing in which a new generation of machine vastly outperforms the world's fastest super-computer, known as "quantum supremacy"
3	politics	The UN's ceasefire in Syria's rebel-held enclave of Eastern Ghouta was cast into doubt less than 24 hours after the Security Council voted to uphold it, as residents woke to regime airstrikes and Iran vowed to carry on fighting in areas it deems held by terrorists.
	sports	Eden Hazard exploded into life and Karim Benzema continued his brilliant scoring run as Real Madrid delivered another goalfest on Saturday in a 4-0 demolition of Eibar.
	business	Wall Street's eternally optimistic forecasters are expecting corporate profit growth to surge by the middle of next year views that are about to collide with reality as hundreds of companies report financial results and update investors on their prospects.
	technology	From ending the opioid epidemic to making fusion power possible, 'Summit' may help researchers meet all sorts of goals. A \$200-million, water-cooled monster that covers an area the size of two tennis courts, the computer, dubbed "Summit," has been clocked at handling 200 quadrillion calculations a second.

Table 20: Example retrieved texts of REGEN over three rounds for AG News dataset.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitation section
- A2. Did you discuss any potential risks of your work?
Ethics Statement section
- A3. Do the abstract and introduction summarize the paper's main claims?
Abstract, Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Yes Section 5

- B1. Did you cite the creators of artifacts you used?
Section 5.1
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Not applicable. Left blank.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Not applicable. Left blank.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
We use public datasets available online, which have also been widely used by other studies.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section 5.1
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Appendix A

C Did you run computational experiments?

Section 5

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Appendix D.3, D.4

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Appendix D.2, table 11. We do not search hyperparameters and use `_one_` hyperparameter set for all tasks.
- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
Yes, section 5.2 (table 1).
- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
Appendix I.1
- D** **Did you use human annotators (e.g., crowdworkers) or research with human participants?**
Section 5.10
- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
Appendix I.2
- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
Appendix I.2
- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
Not applicable. Left blank.
- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
Appendix I.2
- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
Not applicable. Left blank.