

Semi-Supervised Synthetic Data Generation with Fine-Grained Relevance Control for Short Video Search Relevance Modeling

Haoran Li^{1,2}, Zhiming Su², Junyan Yao², Enwei Zhang², Yang Ji², Yan Chen², Kan Zhou², Chao Feng^{2*}, Jiao Ran²

¹National Key Laboratory for Multimedia Information Processing, School of Computer Science, Peking University

²ByteDance Douyin Content Group

haoranli@stu.pku.edu.cn, {sunzhenming, yaojunyan.spacefish, enweizhang, chensong.09, chaofeng.zz, ranjiao}@bytedance.com

Abstract

Synthetic data is widely adopted in embedding models to ensure diversity in training data distributions across dimensions such as difficulty, length, and language. However, existing prompt-based synthesis methods struggle to capture domain-specific data distributions, particularly in data-scarce domains, and often overlook fine-grained relevance diversity. In this paper, we present a Chinese short video dataset with 4-level relevance annotations, filling a critical resource void. Further, we propose a semi-supervised synthetic data pipeline where two collaboratively trained models generate domain-adaptive short video data with controllable relevance labels. Our method enhances relevance-level diversity by synthesizing samples for underrepresented intermediate relevance labels, resulting in a more balanced and semantically rich training data set.

Extensive offline experiments show that the embedding model trained on our synthesized data outperforms those using data generated based on prompting or vanilla supervised fine-tuning(SFT). Moreover, we demonstrate that incorporating more diverse fine-grained relevance levels in training data enhances the model’s sensitivity to subtle semantic distinctions, highlighting the value of fine-grained relevance supervision in embedding learning. In the search enhanced recommendation pipeline of Douyin’s dual-column scenario, through online A/B testing, the proposed model increased click-through rate(CTR) by 1.45%, raised the proportion of Strong Relevance Ratio (SRR) by 4.9%, and improved the Image User Penetration Rate (IUPR) by 0.1054%.

Introduction

Embedding models are fundamental components in many natural language processing applications, including web search, recommendation systems and so on(Huang et al. 2020; Zhao et al. 2023; Xi et al. 2025). Recent studies have explored using LLMs to improve embedding quality (Nie et al. 2024), particularly by synthesizing diverse, high-quality training data to enhance the semantic accuracy of embeddings. Gemini Embedding (Lee et al. 2025) propose a multi-stage prompting strategy to synthesize both retrieval and classification datasets to improve data quality; Qwen3 embedding (Zhang et al. 2025) and KaLM-Embedding (Hu

et al. 2025) leverage Persona Hub (Ge et al. 2024) to guide the data synthesis process, enabling the LLM to generate data from a variety of user personas to construct diverse synthetic train set.

Despite the effectiveness of these methods, existing LLM-based data synthesis methods are predominantly prompt-based and inherently limited by the generative capabilities of the LLM itself. As a result, there remains a noticeable gap between synthetic data and human-curated ground truth data—particularly in low-source domain-specific scenarios such as short video platforms (Zhang et al. 2024). Tang and Yang (2024) proposed FinMTEB, a counterpart to MTEB (Muennighoff et al. 2022; Enevoldsen et al. 2025) that consists of financial domain-specific text datasets, and observed that the state-of-the-art embedding models on MTEB got a significant performance drop on FinMTEB compared to their performance on MTEB. This highlights the necessity of synthesizing data that aligns with domain-specific distributions.

Moreover, current synthesis methods primarily focus on diversity in terms of attributes such as input length, difficulty, and user personas. Little attention has been paid to diversity in terms of semantic relevance, which plays a crucial role in downstream tasks like retrieval and ranking. Most synthesis methods adopt a binary relevance scheme (i.e., relevant vs. irrelevant) and rely on contrastive learning to train embedding models. Although hard negative mining(Zhang et al. 2022) has been widely employed to enhance the model’s sensitivity to subtle relevance differences, the binary formulation introduces a misalignment between the training objective and downstream task requirements. For example, retrieval tasks often demand fine-grained relevance ranking of documents rather than coarse binary relevance classification.

To address these limitations, we introduce a short video relevance dataset, along with two test sets designed for evaluating retrieval and pair classification tasks. All of these datasets are annotated with 4-level fine-grained relevance labels. Furthermore, we propose a **Semi-Supervised Relevance-Aware** data synthesis pipeline, termed **SSRA**. SSRA captures domain-specific data distributions while enabling the controllable generation of samples with predefined semantic relevance levels. It comprises two stages: the first enhances diversity across the synthesized data, and

*Corresponding author.

the second improves the precision of relevance-conditioned generation. By jointly training on annotated domain-specific relevance datasets and high-quality synthetic data generated by SSRA, the embedding model is able to capture fine-grained semantic relevance tailored to the target domain.

We empirically validate SSRA by augmenting training data for the Qwen3-Embedding 4B model. Compared to the baseline trained on annotated data only, our method achieves a 1.73% improvement in nDCG@10 on the retrieval task and a 2.80% gain in average precision (AP) on pair classification. SSRA outperforms prompt-based synthetic data generation and vanilla supervised fine-tuning (SFT) synthetic data generation baselines, demonstrating its effectiveness in synthesizing domain-relevant data with target fine-grained relevance label. Moreover, our analysis highlights the importance of fine-grained semantic relevance diversity in improving downstream performance of embedding models. The embedding model trained with SSRA-augmented data has been deployed in the search enhanced recommendation pipeline of a Douyin’s dual-column scenario, where it achieved significant gains in online A/B testing: CTR +1.45%, SRR +4.9%, and IUPR +0.1054%.

In summary, our contributions are as follows:

- We propose a short video domain relevance dataset and two corresponding test sets for retrieval and pair classification, filling a critical gap in existing benchmarks.
- We introduce **SSRA**, a novel and effective pipeline for synthesizing domain-specific, fine-grained relevance data for embedding model training.
- We empirically demonstrate that increasing the fine-grained semantic relevance diversity of training data improves embedding models’ sensitivity to fine-grained relevance, offering a new perspective on leveraging synthetic data for downstream enhancement.
- Extensive experiments show that embedding models trained on SSRA-generated data achieve substantial improvements in both offline metrics on proposed test sets and real-world online deployment, validating the practical effectiveness of our approach.

Related Work

Embedding Model Training

Embedding models serve as fundamental building blocks for a wide range of natural language processing (NLP) tasks, such as semantic similarity, information retrieval, and classification (Chandrasekaran and Mago 2021; Huang et al. 2020; Keraghel, Morbieu, and Nadif 2024). With the advent of pre-trained language models (PLMs), the learning paradigm of embedding models gradually unify to contrastive learning (Wu et al. 2020; Gao, Yao, and Chen 2021; Izacard et al. 2021). Contrastive learning focuses on constructing positive and negative sample pairs, and optimizes the embedding space by pulling positive pairs closer and pushing negative pairs farther apart. A typical loss function used in this framework is the InfoNCE loss (Oord, Li, and Vinyals 2018). Recent state-of-the-art LLM-based embedding models (Lee et al. 2025; Zhang et al. 2025), whose backbone are

LLMs (Team et al. 2023, 2024; Yang et al. 2025), continue to follow the contrastive learning paradigm, focusing on enhancing model performance through higher-quality positive and negative pairs, as well as improved loss functions based on InfoNCE loss.

Synthetic Data Generation for Embedding Models

Synthetic query generation (Liang et al. 2020; Jeronimo et al. 2023) based on given documents has been widely used to create diverse training data for embedding models. Lee et al. (2024) showed that documents used to synthesize queries may be suboptimal as positive examples. To address this, their method retrieves candidate documents using the generated query and selects the most relevant one via LLM-based scoring. Moreover, many studies directly synthesize positive and hard negative documents as training data for embedding-based retriever tuning. Kim and Baek (2025) leveraged LLMs to synthesize both positive and hard negative documents, aiming to distill the natural language understanding capabilities of LLMs into a lightweight retriever. Shao et al. (2025) proposed a data synthesis pipeline that generates a challenging yet relevant query and a corresponding hard negative document given a document, facilitating the construction of high-quality training triplets.

While these work improved embedding models by synthesizing high-quality and diverse training data, the notion of relevance is still limited to binary labels. This binary supervision introduces a fundamental discrepancy between the training objective of embedding models and the requirements of downstream tasks such as retrieval and reranking. To overcome this limitation, Esfandiarpour et al. (2025) proposed SyCL, a prompt-based approach to generate 4-level relevance documents for a given query using LLMs.

Short Video Relevance Dataset

To facilitate the study of domain-specific synthetic data generation with fine-grained relevance control, we present a novel short video relevance dataset tailored to real-world search and recommendation scenarios with 4-level relevance label. Unlike general-purpose datasets, which often lack domain specificity, our dataset captures the rich diversity and complexity of user intent within the short video domain.

Specifically, we construct a short video search relevance training set along with two evaluation sets, targeting retrieval and pairwise classification tasks respectively. These datasets enables more fine-grained representation learning and evaluation beyond conventional binary relevance setups, and facilitates deeper exploration into semantic matching under domain-specific relevance granularity.

Data Collection

The core of constructing a search relevance dataset lies in building high-quality query-document pairs. However, in the short video domain, each item is inherently multi-modal and lacks explicit textual representation, making it infeasible to directly collect query-document pairs. To address this, we adopt a two-stage data construction pipeline.

Dataset	Size	0	1	2	3	#Queries	#Docs	Avg Query Len.	Avg Doc Len.
Search Relevance Train Set	207,439	65.47	2.85	1.60	30.07	141,950	199,677	7.47	364.69
Retrieval Test Set	10,866	44.17	16.67	14.83	24.33	372	10,605	9.26	454.47
Pair Classification Test Set	3,390	54.45	7.17	1.95	36.43	3,291	3,389	8.14	343.93

Table 1: Statistics and relevance label distributions of the released datasets. Columns named 0–3 indicate the proportion of each relevance label.

Label	Definition
3	Fully and precisely satisfies the user intent with complete semantic relevance to the query.
2	Generally satisfies the user intent and is semantically relevant, but some non-key elements of the content are missing or unmatched.
1	Partially satisfies the user intent; relevance is preserved for key entities or concepts, but significant aspects of the need are unmatched.
0	Neither satisfies the user intent nor maintains any meaningful semantic relevance to the query.

Table 2: Definitions of the 4-level relevance labels used for short video search relevance annotation.

Step 1: Query-Item Collection. We first collect query-item pairs, where each *item* refers to a short video on the short video platform. To obtain high-quality query-item candidates for annotation, we leverage search click logs of Douyin. Specifically, we employ two complementary strategies:

- **Query-driven retrieval:** A set of real user queries is sampled, and candidate items are retrieved using an internal search engine to form query-item pairs.
- **Click-based sampling:** Query-item pairs are directly extracted from search click logs, representing actual user interaction signals.

The training set is constructed by combining data from both strategies to ensure richness and diversity. The retrieval test set exclusively uses query-driven retrieval data to simulate realistic candidate ranking scenarios, while the pairwise classification test set is built solely from click-based query-item pairs.

Step 2: Document Generation. Since short videos lack natural language descriptions, we convert each item into a textual representation using a LLM. Specifically, we utilize Doubao-1.5-Pro-32K¹ to rewrite the item’s available meta-data—namely, its OCR² and ASR³ fields, into coherent textual descriptions. These rewritten descriptions, combined

¹<https://console.volcengine.com>

²Optical Character Recognition (OCR) extracts visual text elements, such as overlaid captions or scene text, from video frames to support semantic interpretation.

³Automatic Speech Recognition (ASR) transcribes spoken audio in the video into text, enabling downstream language understanding tasks.

with the item’s title, form the final *document* associated with each item. The prompt for rewrite OCR and ASR can be found in Appendix-Prompts.

It is worth noting that item ranking in search list is influenced by factors beyond semantic relevance, such as personalization and item popularity. As our objective is to model pure semantic relevance, we filter out query-item pairs with extremely low click-through rates to reduce bias from non-relevance signals.

Data Annotation

After selecting query-document pairs, each query-document pair is annotated with a relevance label on a 4-level scale $\{0, 1, 2, 3\}$ based on a unified annotation guideline. The definitions of each relevance level are detailed in Table 2. This multi-level relevance scheme is specifically designed to reflect nuanced semantic relevance in the context of short video search, capturing domain-specific aspects of content matching and user intent understanding.

To ensure annotation quality and inter-annotator consistency, we adopt a **dual-annotation with adjudication** protocol. Specifically, each pair is independently annotated by two annotators. In cases where their annotations disagree, an expert examiner—well-versed in short video search relevance—is consulted to adjudicate and assign the final label.

Dataset Statistics

The overall statistics of the three datasets are presented in Table 1. It is worth noting that relevance labels 1 and 2 are significantly underrepresented across all datasets, resulting in a highly imbalanced label distribution along the 4-level relevance scale. This observation highlights the necessity of employing data synthesis techniques to enrich mid-level relevance samples, thereby improving the training signal quality and enabling more effective learning of fine-grained relevance distinctions.

Preliminaries

Problem Formulation

Let $\mathcal{D}_{\text{labeled}} = \{(q_i, d_i, s_i)\}_{i=1}^N$ be a labeled dataset, where q_i denotes a query, d_i denotes a document, and $s_i \in \mathcal{S} = \{0, 1, 2, 3\}$ represents the discrete relevance label between the query and document. Additionally, we are given an unlabeled dataset $\mathcal{D}_{\text{unlabeled}} = \{(q_j, d_j)\}_{j=1}^M$ containing a large number of query-document pairs without associated relevance labels. The queries in both datasets are assumed to follow a domain-specific query distribution P_{domain} , i.e., $q_i \sim P_{\text{domain}} \forall q_i \in \mathcal{D}_{\text{label}}, \cup \mathcal{D}_{\text{unlabeled}}$.

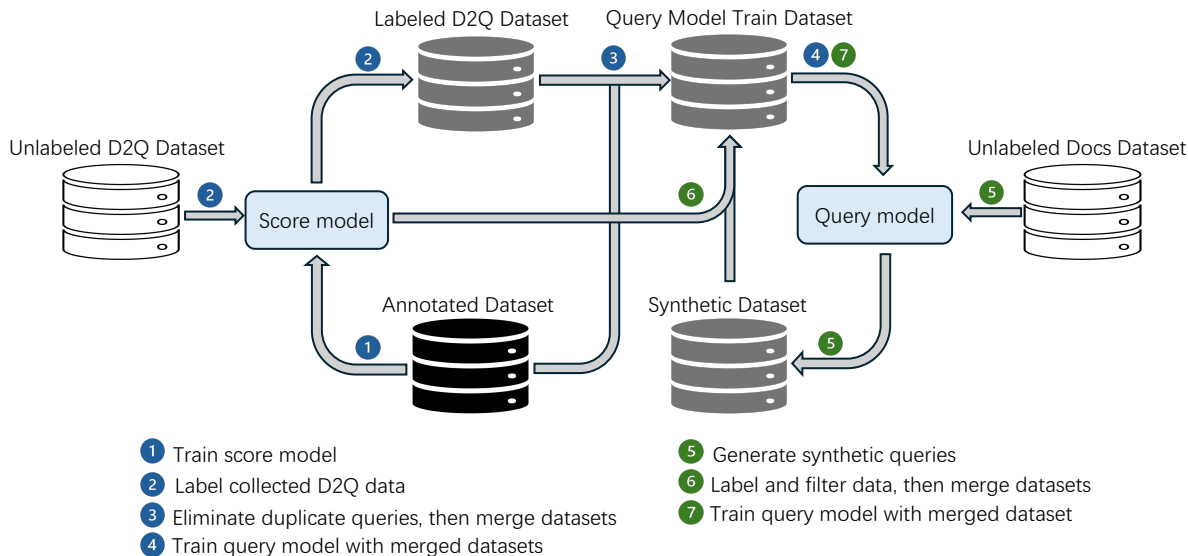


Figure 1: The overview of the proposed two-stage **Semi-Supervised Relevance-Aware** data synthesis (SSRA) pipeline. Steps 1–4 on the left (highlighted in blue) correspond to Stage 1, while steps 5–7 on the right (highlighted in green) represent Stage 2. The term **D2Q** refers to a data structure where each document is associated with multiple queries.

Our goal is to learn a controllable *query model* $f : (d, s) \mapsto \hat{q}$ that, given a document d and a target relevance label $s \in \mathcal{S}$, generates a query \hat{q} such that:

- 1) \hat{q} is sampled from the domain distribution P_{domain_i} , and
- 2) the semantic relevance between \hat{q} and d aligns with the specified score s .

By leveraging this model, we can synthesize query-document pairs with fine-grained, controllable relevance levels that match the target domain distribution.

Method

In this section, we present our proposed SSRA pipeline. The overall framework is illustrated in Figure 1. Our approach is designed as a two-stage process, each targeting distinct objectives to progressively improve query synthesis quality.

Stage 1: Enhancing Query Diversity via Score-Based Re-annotation

As discussed in a previous section, one source of annotated data in our proposed dataset involves retrieving multiple candidate documents for a given query, followed by manual annotation. This data construction paradigm reduces the diversity of generated queries from *query model*, as the *query model* may learn to map multiple distinct documents to the same query during training.

To address this issue, we introduce a relevance classification model, referred to as the *score model*, trained on the labeled dataset to predict discrete 4 relevance labels. Using this model, we collect unlabeled data grouped by document and assign relevance labels to associated queries, enabling the construction of training instances where each document is paired with multiple queries exhibiting different relevance levels. Finally, we merge the resulting synthetic dataset with

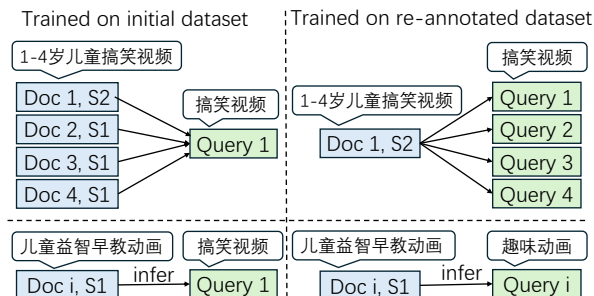


Figure 2: Illustration of our score-based re-annotation strategy in Stage 1. S_i represents the target relevance score. **Left:** In the original labeled dataset, multiple documents are associated with the same high-frequency query, leading to reduced diversity in query generation. **Right:** After Stage 1 re-annotation, each document is linked to multiple queries with varying relevance labels, as labeled by a tuned score model. This enables the query generation model to synthesize diverse queries for unseen documents during inference.

the deduplicated version of the original labeled data to train the initial version of *query model*. This restructured training corpus helps enhance query diversity in the early stage of model learning. Moreover, leveraging real user queries extracted from large-scale unlabeled data enables the model to better approximate the distribution of in-domain queries, thus fostering more realistic and domain-specific query synthesis. Figure 2 shows the change during inference after applying the re-labeled training dataset for *query model*.

Method	0.6B Model					4B Model				
	nDCG@10	AP@≥1	AP@≥2	AP@≥3	Avg	nDCG@10	AP@≥1	AP@≥2	AP@≥3	Avg
Base Model (BM)	71.36	72.63	69.12	66.74	69.50	73.20	72.96	69.14	66.61	69.57
SyCL Modified	71.50	70.34	66.90	64.75	67.33	73.56	71.30	67.43	65.11	67.95
Vanilla SFT	71.88	73.72	70.08	67.38	70.39	74.32	74.22	70.49	67.90	70.87
SSRA	71.97	74.13	70.43	67.80	70.79	74.47	74.92	71.18	68.47	71.52

Table 3: Retrieval and pair classification evaluation results for Qwen3-Embedding of different scales (0.6B vs. 4B). $AP@≥i$ represents Average Precision(AP) computed by treating all samples with relevance labels greater than or equal to i as positive (relevant), and the rest as negative (irrelevant). SSRA consistently improves both nDCG@10 and AP across model sizes.

Stage 2: Enhancing Alignment Between Synthesized Queries and Target Relevance Labels

Despite being trained on carefully constructed data, the query generation model still struggles to reliably synthesize queries that align with specified target relevance labels.

To address this issue while maintaining training efficiency, we adopt an iterative refinement approach that augments the training data with generated queries that better match the desired relevance labels. This is achieved through the following steps:

- **Initial Synthesis:** We first leverage the *query model* trained in Stage 1 to generate a large pool of synthetic queries on the unlabeled document set, conditioned on specific relevance labels.
- **Score-based Filtering:** To ensure the quality of synthesized data, we apply the *score model* tuned on proposed labeled data set to predict the relevance label of each query-document pair. Only those synthetic examples where the predicted label matches the intended target label are retained.
- **Pairwise Consistency Filtering via LLM:** Given that the score model may still introduce errors, we further refine the dataset using a LLM. For each document, we compare pairs of generated queries with differing relevance labels, and filter out cases where the relative order of relevance—judged by the LLM—is inconsistent with the relevance labels. This pairwise filtering step improves the fidelity of relevance differentiation within a single document context.

The remaining high-quality samples are then combined with the original training data from Stage 1 to form an enriched dataset for a second round of query model training. This iterative training process leads to a more robust and relevance-aligned query generator.

Experiments

Experimental Settings

Implementation Details. To train both the *query model* and the *score model*, we adopt Doubao-1.5-Pro-32K as the backbone and both models are trained using a batch size of 16 and an initial learning rate of 1×10^{-5} . To utilize 4-level relevance labels, we utilize infoNCE loss with label as

additional weight of loss⁴.

Following the two-stage SSRA pipeline, we use the final *query model* to synthesize relevance-controlled queries for 1 million distinct short video documents, covering relevance labels of 1, 2, and 3⁵. The synthesized queries are then filtered using the score model to ensure label consistency. These high-quality synthetic queries are combined with the proposed annotated train set to form the final train set for the embedding model.

We employ Qwen3-Embedding⁶ as the backbone for embedding model training, and utilize both in-batch negative sampling and hard negative mining techniques. For We train embedding models at two parameter scales: 0.6B and 4B. For both settings, we adopt the Qwen3-Embedding architecture with lora fine-tuning, following the implementation from the Qwen3-Embedding blog⁷. Specifically, we apply LoRA (Hu et al. 2022) with a rank of 32 and the models are trained with a batch size of 512 and a learning rate of 1×10^{-5} .

More implementation details and corresponding codes are provided in Appendix-Implementation.

Evaluation Settings. We evaluate our methods on the proposed two benchmarks in short video domain: a retrieval test set and a pairwise classification test set. Each benchmark is annotated with four-level relevance labels (0 to 3), capturing fine-grained relevance signals. More details of proposed three data sets can be found in Appendix-Dataset.

For the retrieval test set, we follow the MTEB(Muennighoff et al. 2022) evaluation protocol and report nDCG@10 as the primary metric. For the pairwise classification test set, we also follow the MTEB standard and report Average Precision (AP) at different relevance thresholds (label $≥ 1, 2, \text{ and } 3$) to assess semantic matching quality at varying levels of granularity.

Baselines. We consider several baselines to evaluate the effectiveness of our SSRA pipeline. First, we include a **Base Model** trained directly on the annotated training data without any synthetic augmentation. Second, we evaluate

⁴ $\mathcal{L} = \frac{1}{B} \sum_{i=1}^B s_i \cdot \mathcal{L}_{infoNCE}$, where $s_i \in \{0, 1, 2, 3\}$ is the i -th relevance label in batch, B is batch size.

⁵We use in-batch negative method to get negative pairs, whose relevance label is 0.

⁶<https://huggingface.co/Qwen>

⁷<https://qwenlm.github.io/blog/qwen3-embedding/>

Method	0.6B Model				4B Model			
	nDCG@10	AP@ ≥ 1	AP@ ≥ 2	AP@ ≥ 3	nDCG@10	AP@ ≥ 1	AP@ ≥ 2	AP@ ≥ 3
Binary	71.66	87.36	75.54	54.22	73.73	91.19	77.02	58.39
Multi-Level	71.78	88.89	76.70	52.86	74.23	91.28	77.63	55.25

Table 4: Performance of Qwen3-embedding trained with SSRA augmented data with binary(0,1) and multi-level(0,1,2,3) relevance labels. Note that the AP results are calculated in a balanced-version pair classification test set.

a **Vanilla SFT** baseline, where synthetic data is generated by a query model trained on the annotated dataset and used to augment training dataset. Third, we include a prompt-based baseline, **SyCL Modified**, adapted from SyCL(Esfandiarpour et al. 2025) which originally generates documents of varying relevance labels given a query. We retain its core prompt design but modify the objective to generate relevance-controlled queries given a document, making it a representative of prompt-based query synthesis methods.

Main Results

Table 3 shows our main results on the effectiveness of SSRA-generated, domain-specific synthetic data with controllable fine-grained relevance levels for training embedding models. Compared to the Base Model, our method yields a 1.73% improvement in nDCG@10 on the retrieval task and a 2.80% gain in average precision (AP) on the pair classification task for the Qwen3-Embedding 4B model. For the smaller Qwen3-Embedding 0.6B model, our approach still brings notable gains, with a 0.85% improvement in nDCG@10 and a 1.86% increase in AP.

Our method consistently enhances the performance of Base Model on domain-specific retrieval and pairwise classification tasks. The improvements are observed across models of different parameter scales, demonstrating the robustness of our approach. Furthermore, it outperforms both prompt-based generation methods and direct supervised fine-tuning (SFT) baselines, highlighting the effectiveness of our relevance-controllable query synthesis framework.

Notably, the prompt-based synthesis approach **SyCL Modified** not only fails to improve, but in fact degrades the performance of the embedding model on the pairwise classification task. Specifically, it leads to a drop of 2.17 points in average AP for the 0.6B model and 1.62 points for the 4B model. In contrast, the tuning-based synthesis methods, even when simply trained on annotated data, contributes positively to model performance. Surprisingly, our SSRA method yields substantial gains, improving average AP by 1.29 points on the 0.6B model and 1.95 points on the 4B model. These results highlight the limitations of prompt-based approaches in generating effective domain-specific training data.

Ablation Study and Additional Analysis

We conducted additional experiments, aiming at answering the following key research questions (RQs).

RQ1: Does fine-grained relevance with balanced distribution matters? To assess whether a more diverse relevance distribution (i.e., 4-level relevance labels with a more

balanced label distribution) can lead to improved model performance compared to traditional binary supervision, we conduct controlled experiments using the SSRA-augmented dataset. We evaluate our model using the proposed retrieval test set and a subset of the pair classification task. Since the original pair classification subset contains an imbalanced distribution with only a small proportion of label 1 and 2 samples, we construct a balanced version consisting of 60 test instances for each label (0, 1, 2, 3).

As shown in Table 1, in the original human-annotated dataset, samples with the middle relevance label (12) are underrepresented, and conventional binary annotation schemes tend to label partially matched pairs as non-relevant (i.e., score 0), resulting in a skewed and coarse relevance distribution. In contrast, the SSRA framework, via the *query model*, synthesizes additional middle-relevance samples, leading to a more balanced and fine-grained label distribution.

To simulate the binary relevance setting, we construct a binary subset by filtering out all score-12 samples from the SSRA-augmented dataset. The multi-level subset retains the full range of 4-level relevance labels. We train identical models on both subsets for 2,000 steps under the same training configurations. The final evaluation results are reported in Table 4.

We find that, for the retrieval task, training data with intermediate relevance levels leads to a notable improvement in nDCG@10, highlighting the importance of fine-grained and diverse relevance labels in enhancing retrieval performance.

For the pair classification task, incorporating more training data with intermediate relevance levels can improve the model’s ability to distinguish the middle relevance cases with an improvement on AP@ ≥ 1 and AP@ ≥ 2 . However, this may come at the cost of reduced performance in distinguishing the most relevant pairs from those of lower relevance levels. This finding suggests that while fine-grained relevance annotations are not always equally beneficial across all pair classification thresholds, they remain essential for tasks that require nuanced differentiation, particularly those involving low-relevance or borderline cases.

RQ2: How Do the Two Stages of SSRA Contribute to Retrieval and Pairwise Classification Tasks?

To investigate the contributions of the two-stage SSRA framework across different tasks, we conduct a series of ablation studies, as reported in Table 5. The results reveal that both stages contribute to performance gains in the retrieval task, whereas the improvement in pairwise classification primarily stems from the second stage.

To further analyze the underlying causes of this difference, we examine the impact of each stage on the *query*

Method	0.6B Model		4B Model	
	nDCG@10	AP	nDCG@10	AP
w/o s1&s2	71.44	70.70	74.02	70.34
w/o s2	71.79	70.77	74.23	71.30
SSRA	71.97	74.13	74.47	74.92

Table 5: Ablation results of the two-stage SSRA pipeline using Qwen3-embedding.

Relevance Label	Stage-1 Only	Stage-1 + Stage-2
1	81 / 200	130 / 200
2	80 / 200	131 / 200
3	189 / 200	178 / 200

Table 6: Number of synthesized queries matching the intended relevance labels across training stages.

model from two distinct perspectives:

Effect of Stage 1 on Query Diversity. To quantify the effect of Stage 1 on query diversity, we compare the proportion of duplicate queries generated by two models: (i) a baseline model trained directly on annotated data, and (ii) a model trained with re-labeled data from Stage 1 of SSRA. Specifically, both models generate queries of relevance levels 1, 2, and 3 for 10,000 sampled documents. We then compute the duplicate rate by identifying exact matches among the generated queries.

We observe that the baseline model produces duplicate queries at a rate of 6.57%, whereas the Stage 1-enhanced model reduces this rate to 5.20%, corresponding to a 20.85% relative decrease. This result demonstrates that SSRA Stage 1 effectively improves query diversity. Detailed statistics of query frequencies can be found in the Appendix-Additional Experimental Results.

Effect of Stage 2 on Relevance Alignment. To evaluate whether Stage 2 training enhances the controllability of the query generation model—specifically its ability to synthesize queries that match a target relevance label—we design the following experiment.

We randomly sample 200 documents and use query generation models trained with only Stage 1 and with both Stage 1 and Stage 2 to generate queries intended to match relevance levels 1, 2, and 3, respectively. The generated query-document pairs are then annotated by human annotators using the same 4-level relevance schema. We record the number of generated queries whose annotated relevance label matches the intended label.

As shown in Table 6, the model trained with Stage 2 significantly improves the alignment between intended and actual relevance levels with a 25.43% relative improvement in consistency, demonstrating enhanced controllability and generation precision.

Interpretation of Task-Specific Gains. Based on the above analyses, we interpret the task-specific contributions of SSRA as follows:

For the retrieval task, the model is required to select a subset of relevant documents from a large candidate pool given a query. As such, it is sensitive not only to fine-grained relevance, but also to the diversity of training data. Improved diversity can help mitigate long-tail distribution issues and improve generalization. Therefore, the retrieval task benefits from both stages of SSRA.

In contrast, the pairwise classification task focuses solely on determining the relevance of individual query-document pairs, and does not heavily rely on query diversity. Consequently, it mainly benefits from the second stage, which emphasizes fine-grained relevance supervision.

Metric	CTR	SRR	IUPR
Relative Gain	+1.45%	+4.90%	+0.1054%

Table 7: Relative improvement in online A/B testing metrics after deploying the SSRA-based embedding model.

Online A/B Testing

Based on the embedding model trained with synthetic data from SSRA pipeline, we conducted rigorous online A/B tests in the dual-column scenario of Douyin from May 9 to May 19, 2025, with hundreds of millions of users participating per day. Users were randomly assigned to control and experimental groups via hash bucketing (190 million users each). The embedding model is used to filter out videos irrelevant to the query, ensuring that the displayed videos are closely related to users’ historical search queries.

The results are presented in Table 7, focusing on several key metrics: CTR (Click-Through Rate), SRR (Strong Relevance Ratio), and IUPR (Image User Penetration Rate). SRR is defined as the percentage of content strongly relevant to the user’s search query among the videos ultimately recommended. IUPR refers to the proportion of users who clicked on image content among the total daily active users, which is a core metric in the dual-column scenario.

Conclusion

In this work, we address the limitations of existing LLM-based data synthesis approaches in capturing domain-specific fine-grained semantic relevance. We introduce a relevance dataset in the short video domain and propose **SSRA**, a semi-supervised synthesis pipeline that enables controllable generation of relevance-annotated data aligned with domain distributions. Empirical results demonstrate that SSRA consistently improves embedding model performance on both retrieval and pair classification tasks, outperforming prompt-based and vanilla supervised fine-tuning baselines. Our findings underscore the importance of incorporating fine-grained semantic diversity in synthetic data and point toward promising directions for enhancing embedding models through relevance-aware data generation.

Acknowledgments

We sincerely appreciate our colleagues at ByteDance for their support. They contributed to this work but were not listed as authors. All contributors are listed in alphabetical order by last name: Zirui Guo, Jianheng Ma, Yaling Mou, Lu Ruan, Hailiang Wang.

References

- Chandrasekaran, D.; and Mago, V. 2021. Evolution of semantic similarity—a survey. *Acm Computing Surveys (Csur)*, 54(2): 1–37.
- Enevoldsen, K.; et al. 2025. Mmteb: Massive multilingual text embedding benchmark (2025). *arXiv preprint arXiv:2502.13595*.
- Esfandiarpour, R.; Zerveas, G.; Zhang, R.; Mgonzo, M.; Eickhoff, C.; and Bach, S. H. 2025. Beyond Contrastive Learning: Synthetic Data Enables List-wise Training with Multiple Levels of Relevance. *arXiv preprint arXiv:2503.23239*.
- Gao, T.; Yao, X.; and Chen, D. 2021. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*.
- Ge, T.; Chan, X.; Wang, X.; Yu, D.; Mi, H.; and Yu, D. 2024. Scaling synthetic data creation with 1,000,000,000 personas. *arXiv preprint arXiv:2406.20094*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Hu, X.; Shan, Z.; Zhao, X.; Sun, Z.; Liu, Z.; Li, D.; Ye, S.; Wei, X.; Chen, Q.; Hu, B.; et al. 2025. Kalm-embedding: Superior training data brings a stronger embedding model. *arXiv preprint arXiv:2501.01028*.
- Huang, J.-T.; Sharma, A.; Sun, S.; Xia, L.; Zhang, D.; Pronin, P.; Padmanabhan, J.; Ottaviano, G.; and Yang, L. 2020. Embedding-based retrieval in facebook search. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2553–2561.
- Izacard, G.; Caron, M.; Hosseini, L.; Riedel, S.; Bojanowski, P.; Joulin, A.; and Grave, E. 2021. Unsupervised dense information retrieval with contrastive learning. *arXiv preprint arXiv:2112.09118*.
- Jeronymo, V.; Bonifacio, L.; Abonizio, H.; Fadaee, M.; Lotufo, R.; Zavrel, J.; and Nogueira, R. 2023. Inpars-v2: Large language models as efficient dataset generators for information retrieval. *arXiv preprint arXiv:2301.01820*.
- Keraghel, I.; Morbieu, S.; and Nadif, M. 2024. Beyond words: a comparative analysis of LLM embeddings for effective clustering. In *International Symposium on Intelligent Data Analysis*, 205–216. Springer.
- Kim, M.; and Baek, S. 2025. Syntriever: How to train your retriever with synthetic data from llms. *arXiv preprint arXiv:2502.03824*.
- Lee, J.; Chen, F.; Dua, S.; Cer, D.; Shanbhogue, M.; Naim, I.; Ábrege, G. H.; Li, Z.; Chen, K.; Vera, H. S.; et al. 2025. Gemini embedding: Generalizable embeddings from gemini. *arXiv preprint arXiv:2503.07891*.
- Lee, J.; Dai, Z.; Ren, X.; Chen, B.; Cer, D.; Cole, J. R.; Hui, K.; Boratko, M.; Kapadia, R.; Ding, W.; et al. 2024. Gecko: Versatile text embeddings distilled from large language models. *arXiv preprint arXiv:2403.20327*.
- Liang, D.; Xu, P.; Shakeri, S.; Santos, C. N. d.; Nallapati, R.; Huang, Z.; and Xiang, B. 2020. Embedding-based zero-shot retrieval through query generation. *arXiv preprint arXiv:2009.10270*.
- Muennighoff, N.; Tazi, N.; Magne, L.; and Reimers, N. 2022. Mteb: Massive text embedding benchmark. *arXiv preprint arXiv:2210.07316*.
- Nie, Z.; Feng, Z.; Li, M.; Zhang, C.; Zhang, Y.; Long, D.; and Zhang, R. 2024. When text embedding meets large language model: a comprehensive survey. *arXiv preprint arXiv:2412.09165*.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Shao, R.; Qiao, R.; Kishore, V.; Muennighoff, N.; Lin, X. V.; Rus, D.; Low, B. K. H.; Min, S.; Yih, W.-t.; Koh, P. W.; et al. 2025. ReasonIR: Training Retrievers for Reasoning Tasks. *arXiv preprint arXiv:2504.20595*.
- Tang, Y.; and Yang, Y. 2024. Do we need domain-specific embedding models? An empirical investigation. *arXiv preprint arXiv:2409.18511*.
- Team, G.; Anil, R.; Borgeaud, S.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A. M.; Hauth, A.; Millican, K.; et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Team, G.; Georgiev, P.; Lei, V. I.; Burnell, R.; Bai, L.; Gulati, A.; Tanzer, G.; Vincent, D.; Pan, Z.; Wang, S.; et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Wu, Z.; Wang, S.; Gu, J.; Khabsa, M.; Sun, F.; and Ma, H. 2020. Clear: Contrastive learning for sentence representation. *arXiv preprint arXiv:2012.15466*.
- Xi, R.; Ba, H.; Yuan, H.; Agrawal, R.; Tian, Y.; Kong, R.; and Prakash, A. 2025. Aug2Search: Enhancing Facebook Marketplace Search with LLM-Generated Synthetic Data Augmentation. *arXiv preprint arXiv:2505.16065*.
- Yang, A.; Li, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Gao, C.; Huang, C.; Lv, C.; et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Zhang, X.; Sun, C.; Zhang, R.; and Goh, K.-Y. 2024. The value of AI-generated metadata for UGC platforms: Evidence from a large-scale field experiment. *arXiv preprint arXiv:2412.18337*.
- Zhang, Y.; Li, M.; Long, D.; Zhang, X.; Lin, H.; Yang, B.; Xie, P.; Yang, A.; Liu, D.; Lin, J.; et al. 2025. Qwen3 Embedding: Advancing Text Embedding and Reranking Through Foundation Models. *arXiv preprint arXiv:2506.05176*.

Zhang, Y.; Zhang, R.; Mensah, S.; Liu, X.; and Mao, Y. 2022. Unsupervised sentence representation via contrastive learning with mixing negatives. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 11730–11738.

Zhao, X.; Wang, M.; Zhao, X.; Li, J.; Zhou, S.; Yin, D.; Li, Q.; Tang, J.; and Guo, R. 2023. Embedding in recommender systems: A survey. *arXiv preprint arXiv:2310.18608*.