

CroPS: Improving Dense Retrieval with Cross-Perspective Positive Samples in Short-Video Search

Ao Xie^{*†}, Jiahui Chen[†], Quanzhi Zhu[†], Xiaoze Jiang[‡], Zhiheng Qin, Enyun Yu, Han Li

Kuaishou Technology, Beijing, China

{xieao03, chenjiahui11, zhuquanzhi03, jiangxiaoze, lihan08}@kuaishou.com,

qinzhiheng1991@gmail.com, yuenyun@126.com

Abstract

Dense retrieval has become a foundational paradigm in modern search systems, especially on short-video platforms. However, most industrial systems adopt a self-reinforcing training pipeline that relies on historically exposed user interactions for supervision. This paradigm inevitably leads to a filter bubble effect, where potentially relevant but previously unseen content is excluded from the training signal, biasing the model toward narrow and conservative retrieval. In this paper, we present CroPS (Cross-Perspective Positive Samples), a novel retrieval data engine designed to alleviate this problem by introducing diverse and semantically meaningful positive examples from multiple perspectives. CroPS enhances training with positive signals derived from user query reformulation behavior (query-level), engagement data in recommendation streams (system-level), and world knowledge synthesized by large language models (knowledge-level). To effectively utilize these heterogeneous signals, we introduce a Hierarchical Label Assignment (HLA) strategy and a corresponding H-InfoNCE loss that together enable fine-grained, relevance-aware optimization. Extensive experiments conducted on Kuaishou Search, a large-scale commercial short-video search platform, demonstrate that CroPS significantly outperforms strong baselines both offline and in live A/B tests, achieving superior retrieval performance and reducing query reformulation rates. CroPS is now fully deployed in Kuaishou Search, serving hundreds of millions of users daily.

Introduction

Dense retrieval has emerged as a foundational paradigm in modern search systems due to its efficiency and effectiveness in matching semantically rich queries to relevant documents (Zhang et al. 2022; Ma et al. 2024). In short-video search platforms, dual-encoder retrievers (often referred to as “two-tower” models) are widely adopted, where separate encoders are used to independently embed user queries and candidate videos into a shared latent space. Typically, these retrievers are trained using a data-driven and self-reinforcing paradigm: as shown in Figure 1 (gray block), historical user

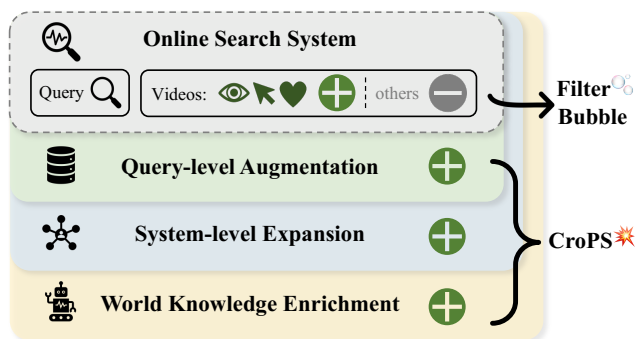


Figure 1: Overview of CroPS, where “+” (“-”) represents positives (negatives). It introduces positives from three complementary sources: Query-level Augmentation, System-level Expansion, and World Knowledge Enrichment.

interactions (*e.g.*, watching, clicking, or liking) from a production search system serve as supervision. Videos previously displayed for a given query are labeled as positive examples, while negative samples are drawn from content that was filtered out during earlier retrieval stages or never exposed to the user (Zheng et al. 2024). While this paradigm facilitates iterative model improvement through continuous data accumulation, it also poses a critical challenge: the filter bubble effect. Because only historically exposed content is considered as positive supervision, potentially relevant but previously unseen videos are excluded from the positive set, even incorrectly labeled as negatives. This bias restricts the model’s ability to surface novel or diverse content, leading to conservative and narrow retrieval behavior. The resulting lack of diversity impairs user experience and limits the system’s capacity to support exploratory search.

Previous research has primarily focused on improving retrieval performance through architectural innovations (Zhang et al. 2022; Khattab and Zaharia 2020) or refined negative sampling strategies (Karpukhin et al. 2020; Yang et al. 2024). However, these efforts remain constrained by the limitations of the self-reinforcing training pipeline and thus cannot fully escape the filter bubble (Meghwani et al. 2025; Ren et al. 2021). In contrast, we propose that positive

^{*}Contribution during internship at Kuaishou Technology.

[†]These authors contributed equally.

[‡]Corresponding author.

sample enrichment offers a powerful and complementary solution that has been largely overlooked. By introducing relevant content beyond the scope of historical exposure, we can effectively break through the boundaries of existing data.

In this paper, we propose CroPS (Cross-Perspective Positive Samples), a retrieval enhancement data engine that incorporates diverse positive signals from multiple perspectives to break the limitations imposed by the filter bubble (see Figure 1). CroPS enriches the training signal with previously unobserved yet semantically relevant examples from three complementary sources. First, it captures **query-level augmentation** by identifying the user query reformulation behavior. When users issue a follow-up query with similar semantics after an unsatisfactory search, CroPS treats videos consumed after such reformulations as additional positives for the original query. This leverages intent continuity to expand the positive set beyond what the system originally retrieved. Second, CroPS introduces **system-level augmentation** by bridging search and recommendation systems. It incorporates videos from the recommendation flow that align semantically with the user’s search query, thus breaking the silo between independently functioning subsystems and leveraging user engagement data across systems. Finally, CroPS performs **world knowledge augmentation** by using large language models as virtual retrievers. These models, endowed with rich semantic and factual knowledge, are used to synthesize high-quality training signals that reflect relevant content beyond what is observable in user interaction logs. Meanwhile, it simulates user behavior where, upon failing to obtain satisfactory information within the current app, users seek knowledge from alternative sources. This is regarded as a form of cross-platform knowledge integration, which can effectively mitigate the effects of information cocoons. Together, these augmentations significantly expand the positive space, allowing the retriever to learn a more comprehensive representation of relevance and generalize beyond historically exposed content.

However, naively treating all these heterogeneous positives as equal during training leads to suboptimal learning. We propose a Hierarchical Label Assignment (HLA) strategy that assigns different levels to positive samples based on their origin and reliability. This hierarchical supervision allows the model to learn fine-grained ranking behaviors aligned with the system objectives. For instance, by assigning higher labels to query-level augment positives, the model learns to generalize across reformulated queries, potentially reducing the frequency of user query rewrites. To support HLA-based training, we introduce H-InfoNCE, a novel loss function tailored to hierarchical supervision, enabling more targeted optimization across different levels of relevance. Unlike InfoNCE (Oord, Li, and Vinyals 2018), which individually compares each positive sample with its corresponding negatives, H-InfoNCE requires that all positive samples from higher tiers be jointly contrasted with negatives from lower tiers in a single step. By leveraging efficient and large-scale positive-negative contrastive learning, H-InfoNCE empowers our model to achieve rapid and effective learning, leading to notable performance gains and providing meaningful insights for practical industrial applications.

In summary, this work makes three key contributions. (1) We identify a fundamental limitation in industrial dense retrieval systems: the presence of filter bubbles caused by self-reinforcing training paradigms, and argue that introducing positive samples beyond the historical exposure space is an effective and underexplored solution. (2) We propose CroPS, a novel framework that enriches the positive training set by incorporating semantically relevant examples from query reformulations, cross-system user feedback, and external world knowledge via large language models. To support learning from this heterogeneous supervision, we introduce a Hierarchical Label Assignment (HLA) strategy along with a tailored H-InfoNCE loss, enabling the model to capture fine-grained ranking preferences aligned with system-level goals. (3) We conduct extensive offline and online evaluations on Kuaishou Search, a large-scale short-video search platform. The results demonstrate that CroPS not only achieves state-of-the-art retrieval performance in offline experiments, but also significantly improves user engagement and reduces query reformulation rates in online A/B testing. CroPS has been successfully deployed in Kuaishou Search and now serves as a core component of the search infrastructure for hundreds of millions of active users.

Related Works

Dense Retrieval. The modern model commonly employs a dual-encoder architecture, which is well adapted to approximate nearest neighbor (ANN) (Johnson, Douze, and Jegou 2021) retrieval. Methods in dense retrieval can generally be categorized into five main areas: retrieval-oriented continued pre-training (Su et al. 2024; Oguz et al. 2022; Chang et al. 2020), negative sampling strategies (Meghwani et al. 2025; Yang et al. 2024; Zhou et al. 2022), structural augmentation techniques (Luan et al. 2021; Khattab and Zaharia 2020; Humeau et al. 2020), the exploration of generative paradigms (Li et al. 2025; Li, Zhou, and Dou 2024; Tang et al. 2024; Wang et al. 2025), and the applications in specific industrial scenarios (Lin et al. 2024; Rossi et al. 2024; Zheng et al. 2022). These works have not conducted an in-depth study of positive samples and failed to address the information cocoon effect in the retrieval domain. Furthermore, the reliance on complex post-interaction models like the Poly-encoder (Humeau et al. 2020) presents significant challenges for integration with ANN-based systems, which in turn hampers practical deployment. Unlike previous works, our architecture-agnostic approach can be applied to enhance any model using CroPS without introducing any inference overhead, making it well-suited for online applications. At the same time, it mitigates the information cocoon effect in retrieval systems by leveraging positive sample enhancement and a hierarchically stratified contrastive loss, thereby increasing the diversity of search results and improving the overall user experience.

Contrastive Learning. Contrastive learning has become a fundamental paradigm for representation learning in both vision and language domains (Wu et al. 2018; He et al. 2020; Radford et al. 2021; Jiang et al. 2022). In retrieval, contrastive objectives such as InfoNCE (Oord, Li, and Vinyals 2018) and its variants (Li et al. 2021; Magnani et al.

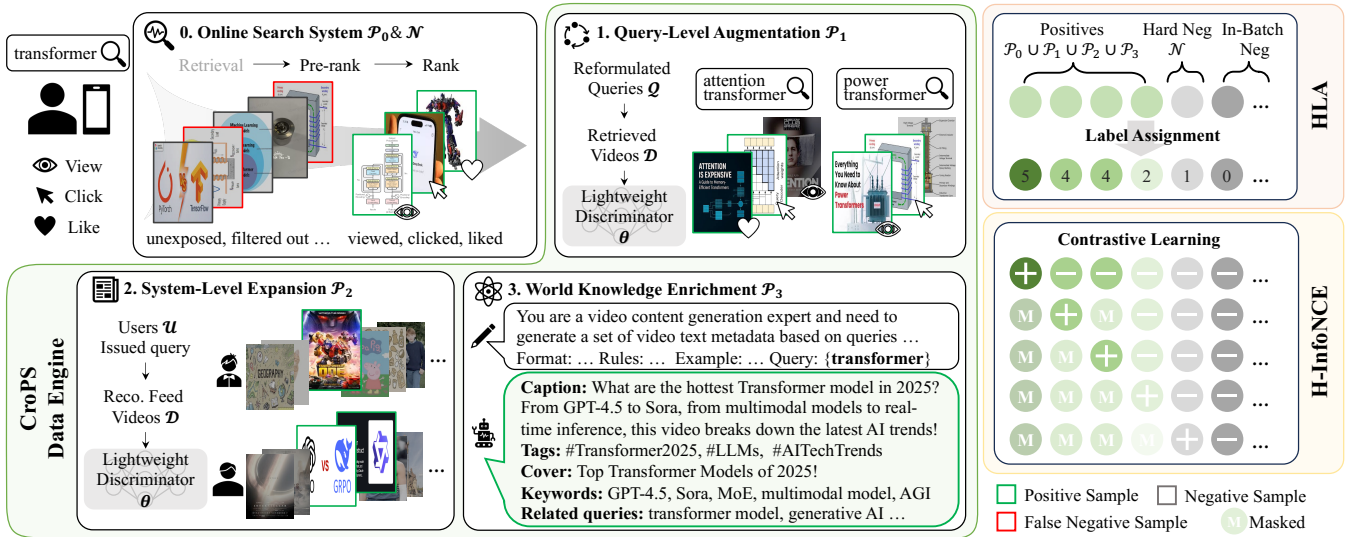


Figure 2: The overall framework of the CroPS, where “+” (“-”) represents positives (negatives). It consists of three components: CroPS Data Engine, Hierarchical Label Assignment (HLA), and H-InfoNCE Optimization. CroPS Data Engine expands the positive samples from: (1) Query-level Augmentation, (2) System-level Expansion, and (3) World Knowledge Enrichment.

2022; Hoffmann et al. 2022) form the backbone of training strategies by maximizing the similarity between query-positive pairs while pushing apart query-negative pairs. Various works have extended this framework to enhance efficiency and robustness, including hard negative mining (Yang et al. 2024), in-batch negatives (Radford et al. 2021), and list-wise formulations (Hoffmann et al. 2022). However, most of these approaches assume a flat binary relevance structure and fail to capture the nuanced semantic relationships between different types of positive and negative samples. Recent efforts have explored hierarchical or weighted contrastive learning, yet these typically rely on pre-defined heuristics or coarse metadata for label assignment. In contrast, our proposed HLA (Hierarchical Label Assignment) introduces a data-driven, semantically grounded label hierarchy over both positive and negative samples. Built upon this, we propose H-InfoNCE, a stratified contrastive loss that enables the model to learn graded relevance and fine-grained ranking preferences. It better reflects real-world retrieval scenarios, where training data often contain heterogeneous supervision signals with varying confidence levels.

Methodology

In industrial short-video search and recommendation, textual information—comprising user-edited titles, captions, keywords, and OCR-extracted text fields—offers an effective, cost-efficient video representation (Zheng et al. 2024, 2022). As a standard practice, our retrieval model also utilizes this textual information as input to encode videos.

CroPS Data Engine

Different from prior works that primarily focus on better negative sampling strategies for retriever training (Meghani et al. 2025; Yang et al. 2024), CroPS takes a com-

plementary perspective by enriching the positive sample set. By incorporating multiple sources of semantically relevant items, it breaks through the limitations of the conventional feedback loop and mitigates the filter bubble effect.

As shown in Figure 2, given a query q , the CroPS data engine collects the positive set $\mathcal{P} = \mathcal{P}_0 \cup \mathcal{P}_1 \cup \mathcal{P}_2 \cup \mathcal{P}_3 = \{d_1^+, d_1^+, d_2^+, \dots, d_k^+\}$ from the four perspectives below, and the negative set $\mathcal{N} = \{d_{k+1}^-, d_{k+2}^-, \dots, d_n^-\}$.

Online Search System. The retrieval process in short-video search systems typically follows a multi-stage pipeline consisting of recall, pre-ranking, and ranking, which sequentially filters and ranks candidate videos. Following common practice (Zheng et al. 2024), we treat videos that were clicked or liked during the ranking stage as a basic positive set \mathcal{P}_0 . For the negative sample set \mathcal{N} , we sample hard negatives from two sources: videos that were not exposed to users in the ranking stage, and those that filtered out between the pre-ranking and ranking stages. This sampling strategy strikes a balance between positive relevance and negative difficulty, ensuring a meaningful training signal. However, relying solely on interaction-based labels can reinforce filter bubble effects: relevant but previously unexposed content is systematically excluded from positive training, introducing bias and limiting the model’s ability to surface novel but pertinent results. As illustrated in Figure 2, for the query “transformer”, user interactions are dominated by videos about transformer models in deep learning or the Transformers franchise, while videos related to electrical transformers are rarely exposed and thus incorrectly categorized as negatives, despite being semantically relevant. To mitigate such false negatives, CroPS enriches the positive set with diverse and semantically relevant samples, promoting broader content coverage and improving retrieval robustness.

Query-Level Positive Augmentation. Query reformulation

is a common user behavior in search systems, often indicating a continued effort to refine or clarify an underlying information need. When a user issues a semantically related follow-up query shortly after an initial one, the videos consumed in response to the new query are likely to remain relevant to the original intent, even if they were not retrieved in the initial search. These samples typically lie beyond the retrieval scope of the original query and thus offer valuable signals for expanding the positive set, enhancing both recall diversity and model generalization. To identify such samples, we employ a lightweight semantic discriminator $\theta(\cdot)$ to measure the relevance between the original query q and candidate videos d_{ij} consumed after reformulated queries. The query-level positive set is defined as:

$$\mathcal{P}_1 = \bigcup_{q_i \in \mathcal{Q}} \{d_{ij} \in \mathcal{D}_i \mid \theta(q, d_{ij}) > \alpha\} \quad (1)$$

Here, q is the original query, \mathcal{Q} is the set of reformulated queries issued by the same user within a short time window (e.g., 90 seconds), \mathcal{D}_i is the set of videos the user interacted with (e.g., clicked or watched) after issuing $q_i \in \mathcal{Q}$, and α is a predefined semantic relevance threshold. As illustrated in Figure 2, consider a user initially searching for “transformer” but receiving unsatisfactory results. They may reformulate the query to “power transformer” to clarify the intent. The relevant videos retrieved and consumed under the reformulated query can serve as positive samples for the original query “transformer”. This approach not only corrects potential false negatives in \mathcal{N} but also reduces the overemphasis on dominant interpretations, thus promoting more balanced and accurate retrieval.

System-Level Positive Expansion. In addition to the search system, short-video platforms operate a recommendation system that independently serves a large portion of user traffic. The recommendation feed typically generates extensive user interaction data, which contains valuable signals that can help mitigate the information cocoon effect. CroPS leverages this opportunity by bridging the gap between search and recommendation, enabling cross-system positive signal integration to further enhance training supervision. Specifically, for a given query q , we identify a set of users \mathcal{U} who have issued this query and retrieve up to 100 videos \mathcal{D}_i that each user $u_i \in \mathcal{U}$ interacted with from the recommendation feed, within a temporal window centered around the query timestamp. A lightweight semantic discriminator $\theta(\cdot)$ is employed to evaluate the relevance between the query q and each candidate video d_{ij} . Those with scores exceeding a threshold α are included in the system-level positive set, formally defined as:

$$\mathcal{P}_2 = \bigcup_{u_i \in \mathcal{U}} \{d_{ij} \in \mathcal{D}_i \mid \theta(q, d_{ij}) > \alpha\} \quad (2)$$

This strategy enables the model to incorporate semantically relevant signals that are not captured by the search pipeline, thereby increasing the coverage of positive examples and mitigating exposure bias. Compared to search-derived samples, recommendation-based positives are typically more timely and better aligned with users’ personal interests, offering a valuable supplement to query-based supervision.


	Sample Types	Label
Positive  Negative	Query-level aug. positives	5
	System-level aug. positives	4
	World-knowledge aug. positives	4
	Clicked videos in rank	4
	Exposed videos in rank	3
	Unexposed videos in rank	2
	Filtered between pre-rank and rank	1
In-batch negatives	0	

Table 1: Hierarchical Label Assignment (HLA) for positive and negative samples.

World Knowledge Enrichment via LLMs. Despite the incorporation of query-level and system-level positive sample enrichment, the filter bubble effect may still persist due to the inherent limitations of existing search and recommendation engines, as well as the finite scope of the platform’s video inventory. To further mitigate this constraint and break through the platform’s content exposure boundaries, we introduce a world knowledge enrichment strategy that leverages large language models (LLMs) as an external knowledge source. Specifically, the LLM is employed as a pseudo-retriever to synthesize high-quality positive examples \mathcal{P}_3 . As shown in Figure 2, we adopt a one-shot prompting strategy: the LLM is provided with a query and an exemplar video deemed relevant, and is then instructed to generate other video descriptions aligned with the query. These synthetic examples serve as positive samples that enrich the training set with diverse linguistic expressions, conceptual variations, and external knowledge not represented in the platform’s existing content or logs. Importantly, these generations capture latent semantic associations and factual information from the broader world, enabling the retrieval model to generalize beyond the scope of historical user interactions. By integrating LLM-generated positives into the training process, we expand the model’s relevance coverage, enhance its capacity to surface novel but meaningful content, and further reduce the impact of content exposure bias.

Hierarchical Label Assignment

The CroPS data engine introduces positive samples from multiple perspectives. However, these samples vary in importance and reliability. Treating them uniformly can lead to suboptimal learning outcomes. To address this, we propose a Hierarchical Label Assignment (HLA) mechanism. Rather than assigning equal importance to all positive examples, HLA allocates discrete label levels ranging from 0 to 5¹ to reflect varying degrees of relevance and to guide the retriever’s ranking preferences accordingly. This hierarchical supervision enables the model to learn fine-grained semantic distinctions and supports flexible optimization for different training objectives.

Formally, for a query q and its sample set $\mathcal{S} = \mathcal{P} \cup \mathcal{N} = \{d_1, d_2, \dots, d_n\}$ that contains n positive/negative videos,

¹The label range is flexible; only the relative order is essential.

HLA assigns n labels $\{l_1, l_2, \dots, l_n\}$ to each sample d_i based on its origin, as detailed in Table 1. Query-level augmented positives receive the highest label (5) because they most directly reflect refined user intent: the reformulated queries are typically issued by users when initial results are unsatisfactory, and the subsequent interactions often provide highly relevant signals. A label of 4 is given to system-level and world-knowledge augmented positives, as well as clicked videos, all of which provide strong relevance signals from different perspectives. Exposed but unclicked videos in the ranking stage are labeled 3, indicating moderate relevance. Unexposed videos in the ranker’s output receive label 2, capturing weak or uncertain relevance. Videos filtered between pre-rank and rank are assigned label 1, while in-batch negatives receive the lowest label of 0. This hierarchical scheme enables the retriever to exploit richer supervision beyond binary feedback and learn to rank content in accordance with nuanced relevance signals, ultimately improving both recall diversity and user satisfaction.

H-InfoNCE Optimization

To effectively leverage the multi-grade supervision provided by HLA, we adopt the Hierarchical InfoNCE (H-InfoNCE) loss function. Unlike the standard InfoNCE loss, which assumes binary relevance, H-InfoNCE imposes a label-aware contrastive structure: for a given positive sample with label l , only samples with strictly lower labels ($< l$) are considered as negatives. This formulation, shown in Eq. 3, better aligns with the graded supervision and guides the model to learn fine-grained, list-wise ranking preferences.

$$\mathcal{L} = - \sum_{d_i \in \mathcal{S}} \log \frac{\exp\left(\frac{\text{sim}(q, d_i)}{\tau}\right)}{\sum_{d_j \in \{d_i\} \cup \{d_k \in \mathcal{S} | l_i > l_k\}} \exp\left(\frac{\text{sim}(q, d_j)}{\tau}\right)} \quad (3)$$

We ensure training scalability by efficiently implementing H-InfoNCE, using a mask matrix to filter non-comparable samples by label (see Figure 2, where samples with labels $\geq l$ are masked) and a label-indexed data structure to organize inputs. This allows computation of all hierarchical contrastive losses in one forward pass, matching standard InfoNCE speed and making it suitable for industrial-scale deployment for large-scale data.

Experiments

Dataset, Metrics and Details

Dataset. Public datasets for information retrieval evaluation often fail to meet the demands of real-world applications. These datasets only record single search queries, ignoring the complex user behavior patterns found in actual systems. Thus, we create a new dataset, called CPSQA, that captures diverse user behaviors, including query reformulation sequences and cross-domain consumption patterns covering search and recommendation systems. CPSQA is developed using real user interaction logs from Kuaishou Search, capturing authentic user behavior patterns at an industrial scale. Specifically, we utilize logs from June 11, 2025, for training (500 million samples) and logs from June 13, 2025, for testing (10,000 samples). The CPSQA reflects the scale of

online production environments, incorporating a candidate item pool of about 1 billion items.

Metrics. For evaluation, we employ Recall@100 on two test sets: *click-through data* (CT) and *query reformulation data* (QR). The CT test set treats clicked videos as positives, demonstrating immediate user preferences. The QR test set comprises videos consumed following query reformulation, which are validated by a semantic discriminator based on a predefined relevance threshold. It is designed to evaluate the model’s ability to generalize beyond historical exposure patterns and mitigate the filter bubble effect. In addition, we conduct NDCG@4 evaluation on the manually annotated test dataset. NDCG@4 is utilized to assess the recall model’s ability to rank the top 4 results, serving only as a reference monitoring metric. The primary evaluation metric for the recall model is still the recall rate.

Implementation Details. The CroPS data engine comprises three key components. For query-level positive sample augmentation, we identify user query reformulation behavior by detecting pairs of semantically similar queries issued by the same user within a 90-second window. The lightweight discriminator $\theta(\cdot)$ that we used is a pre-trained 6-layer Transformer model. The similarity threshold α in Eq. 1 is empirically set to 0.6. The system-level positive sample expansion module follows the same design. For world knowledge enrichment, we utilize the Qwen2.5-14B (Qwen et al. 2025) as a pseudo-retriever to synthesize external positive examples. We generate 35 million synthetic positive samples. During Training, all retrieval models, including our proposed method and the baselines, adopt a dual-encoder architecture. We initialize both the query and document encoders using the Qwen2.5-0.5B pre-trained language model. The query encoder takes the raw query text as input, while the document encoder is fed with the textual content of the photo. The maximum input sequence length is limited to 128 tokens. We set the temperature factor τ in Eq. 3 as a learnable parameter with a initial value of 0.05. We train all models using the Adam optimizer (Kingma and Ba 2015), with a learning rate of 2e-5 and a weight decay of 1e-4.

Main Results

To comprehensively evaluate the effectiveness of CroPS, we conduct experiments on the CPSQA dataset and report three key metrics: NDCG@4, and Recall@100 on CT and QR.

The comparison methods fall into three major categories. Classical methods: BM25 (Robertson and Zaragoza 2009) as a probabilistic ranking baseline and NCE (Gutmann and Hyvärinen 2010) as a traditional contrastive learning approach. Neural network methods: DPR (Karpukhin et al. 2020) with dual-encoders, ANCE (Xiong et al. 2021) using dynamic hard negative sampling, and ADORE+STAR (Zhan et al. 2021) for stable optimization. Negative sampling strategies: TriSampler (Yang et al. 2024) with principled sampling and FS-LR (Zheng et al. 2024), which introduces multi-level negative labels.

As shown in Table 2, CroPS demonstrates significant advantages in retrieval performance that directly impact user experience. On the CT dataset, CroPS achieves a Recall@100 of 69.1%, surpassing the best baseline (TriSam-

Method	Recall@100 \uparrow		NDCG@4 \uparrow (%)
	CT(%)	QR(%)	
BM25	42.9	22.5	64.8
NCE	53.6	27.5	65.4
DPR	56.0	30.7	66.5
ANCE	56.9	31.3	67.1
ADORE + STAR	59.4	31.9	67.4
TriSampler	59.8	32.2	66.9
FS-LR	59.6	33.0	66.0
CroPS (Ours)	69.1	40.1	67.0

Table 2: Performance comparison of different methods on search ranking tasks.

pler) by 9.3%, demonstrating advanced capability in obtaining content that corresponds with user search intent. More importantly, CroPS obtains 40.1% Recall@100 on QR, which significantly outperforms existing models. It means that the users can find relevant content in their first search attempt, reducing query refinement needs and improving search satisfaction. This demonstrates CroPS’s ability to break through information silos and understand user search intent even with imprecise initial queries. While NDCG@4 measures ranking quality based on human-annotated relevance judgments, CroPS achieves competitive performance at 67.0%, remaining on par with top-performing baselines. The key advantage lies in combining extraordinary recall capabilities with maintained ranking precision. This means that with comparable ranking quality, users can immediately find exactly what they’re looking for without multiple search iterations. The results establish CroPS as setting new standards for industrial search by delivering both comprehensive content coverage and superior user experience through its multi-source positive sample strategy.

Ablation Study

We conduct ablation studies focusing on core components of CroPS: Data Engine, HLA, and H-InfoNCE Optimization.

Effectiveness of the CroPS Data Engine. We progressively incorporate three augmentation strategies—query-level augmentation, system-level augmentation, and world knowledge enrichment—on top of the baseline that only uses online search data. As shown in Table 3, each augmentation consistently improves performance. Adding query-level augmentation improves Recall@100 by +3.7% on CT and notably +3.1% on QR, showing the value of capturing user intent continuity through query reformulation behavior, with the QR improvement being the largest single-strategy gain across all individual augmentations. In addition, including system-level enhancement brings a greater gain of +2.4% in CT and +1.3% in QR, confirming the benefit of leveraging recommendation signals to enrich exposure diversity. Incorporating world knowledge augmentation using LLMs results in the further gain, achieving a final Recall@100 of 69.1% (CT) and 40.1% (QR), as well as the best NDCG@4 score of 67.0%. This shows that LLM-generated examples com-

Model	Recall@100 \uparrow		NDCG@4 \uparrow (%)
	CT(%)	QR(%)	
(1) <i>CroPS Data Engine Ablation</i>			
Baseline	59.6	33.0	66.0
+ Query-level Aug.	63.3	36.1	66.5
+ System-level Aug.	65.7	37.4	66.7
+ World Knowledge Aug.	69.1	40.1	67.0
(2) <i>Hierarchical Label Assignment Ablation</i>			
CroPS [†] (binary label)	59.9	32.1	66.7
CroPS [‡] (query-level aug. label = 4)	67.1	38.4	66.8

Table 3: Data augmentation and HLA ablations.

Loss	Recall@100 \uparrow		NDCG@4 \uparrow (%)	Speed \downarrow
	CT(%)	QR(%)		
InfoNCE	67.8	38.9	66.9	178h
Softmax-CE	65.3	37.8	66.8	89h
H-InfoNCE	69.1	40.1	67.0	88h

Table 4: Contrastive optimization comparison.

plement the in-platform data by injecting external semantics and broader world associations. These results validate that each augmentation source addresses a unique dimension of the filter bubble, and their combination yields additive gains in both recall and ranking performance.

Effectiveness of Hierarchical Label Assignment (HLA).

To evaluate the impact of our hierarchical labeling strategy, we conduct two ablation studies as shown in the second part of Table 3. First, we compare the CroPS model with a variant CroPS[†], where hierarchical labels are simplified into binary form: samples with label ≥ 4 are treated as positives, and others as negatives. This variant uses a standard InfoNCE loss for optimization. The results show a significant performance drop, with Recall@100 decreasing by 9.2% on CT and 8.0% on QR compared to the CroPS. This highlights the importance of preserving label granularity, as HLA provides more informative supervision that helps the model distinguish varying degrees of relevance. We further investigate the effect of assigning the highest label (*i.e.*, 5) to the query-level augmented positives. As discussed in HLA, query-level augmented positives receive the highest label because they represent the most authentic manifestation of user preferences, as these reformulation behaviors inherently encode what users truly seek after experiencing initial search outcomes. To validate this design choice, we modify this setting in CroPS[‡] by assigning a lower label of 4 to such positives. This change leads to noticeable degradation on Recall@100 by 2.0% on CT and 1.7% on QR, confirming that query-level reformulation signals are indeed particularly strong indicators of retrieval relevance and should be weighted accordingly in training. Together, these results demonstrate that HLA offers finer-grained supervision, which significantly contributes to CroPS’s retrieval effectiveness.

Model Type	Model	Ratio _{rank} ↑	Ratio _{show} ↑
Dense Model	Baseline	29.8%	32.5%
	CroPS	40.9%	44.3%
Sparse Model	Baseline	20.8%	29.1%
	CroPS	33.8%	44.7%

Table 5: The improvements of CroPS method in Ratio_{rank} and Ratio_{show} compared to the production baseline.

Model Type	CTR ↑	LTR ↑	RQR ↓
Dense Model	+0.869%	+0.483%	-0.646%
Sparse Model	+0.783%	+0.423%	-0.614%

Table 6: The improvements of CroPS in online A/B test.

Efficiency of H-InfoNCE. We compare H-InfoNCE loss against standard InfoNCE and Softmax cross-entropy (Magnani et al. 2022) losses. The results are shown in Table 4. InfoNCE compares each positive sample with its corresponding negative instances within a single training step. It leads to the slowest training efficiency. Although Softmax-CE trains faster, it yields the worst results. The high-level samples are mistakenly introduced when computing the low-level contrastive loss, even with score regularization to address this issue in Softmax-CE. In contrast, H-InfoNCE not only achieves the best performance across all metrics, but also maintains high training efficiency. By computing multi-level contrastive loss in a single forward pass, it significantly improves both accuracy and efficiency. Such efficiency is particularly beneficial in industrial scenarios, where models must be trained frequently on large-scale data with limited computational resources.

Online Testing

We validate the proposed CroPS method with online A/B testing on Kuaishou Search.

Scalability. Our CroPS is agnostic to the model architecture. We evaluate our method on both dense and sparse model. The dense model is the previously mentioned Qwen2.5-0.5B, which predominantly processes text input. The sparse model employs the MLP structure and utilizes ID features as input. Consistent improvements are observed in Table 5 and Table 6, when applying CroPS to these two distinct model types, demonstrating the scalability of our approach.

Recall Capability. We investigate how our recall results are exposed in the ranking stage. It includes two metrics: Ratio_{rank} and Ratio_{show}. Ratio_{rank} represents the ratio of our recall set within the ranking set, while Ratio_{show} indicates the ratio of our recall set within the show set (showing to users). As shown in Table 5, the CroPS on dense model contributes to a 11.1% increase in Ratio_{rank}, a 11.8% increase in Ratio_{show}. The higher recall ratio reflects better retrieval of relevant content.

User Satisfaction. To further evaluate the user’s interaction with the search results page after searching, we introduce the metrics of CTR, LPR (long-play rate), and RQR

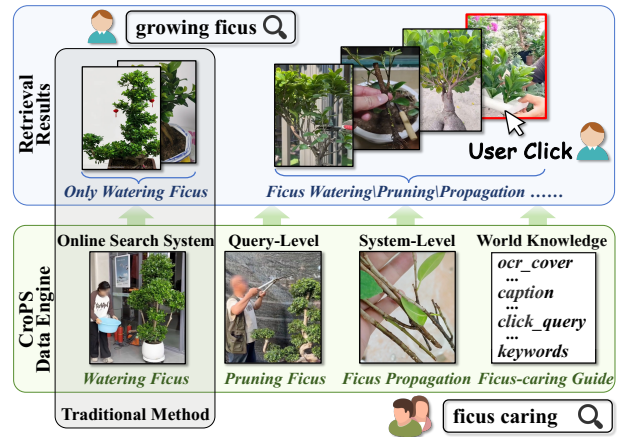


Figure 3: Case Study. Blue block indicates test results, while green block shows evidence traces from our method.

(reformulated-query rate). CTR assesses whether the search results meet user needs and drive click behavior, LPR measures the proportion of users who watch a video for an extended duration, and RQR is the proportion of users who change their query and initiate a search again. The results of CroPS are shown in Table 6, our comprehensive CroPS method on dense model contributes to a 0.869% increase in CTR, a 0.483% increase in LPR and a 0.646% decrease in the RQR. The observed high click-through and long watch times, coupled with a low rate of query reformulation, indicate that users successfully located their intended content. This demonstrates that our model effectively captures and fulfills user search intent.

Case Study

Considering query “growing ficus”, as shown in the blue block of Figure 3, traditional methods retrieve only watering or fertilizing videos, while CroPS retrieves diverse content including propagation, pruning, and disease prevention videos. The user ultimately clicks on the propagation video from our results. As shown in the green block of Figure 3, we further explore the supporting evidence of the CroPS data engine by retrieving similar queries from the training data using the original query as an anchor. We observe that traditional click-based training data primarily consists of basic care videos. In contrast, CroPS’s data engine enriches training samples with diverse content like propagation techniques, pruning guides, and pest control from multiple channels, capturing various user intents beyond clicked results. This design breaks information bubbles by expanding training diversity, better satisfying varied user needs.

Conclusion

To address information cocoons in industrial dense retrieval systems, we propose CroPS, which enriches the positive training signal with multiple perspectives, offering relevant yet novel supervision. Hierarchical Label Assignment and H-InfoNCE enhance fine-grained retrieval semantics. Future work will integrate CroPS with generative retrieval methods.

References

- Chang, W.-C.; Felix, X. Y.; Chang, Y.-W.; Yang, Y.; and Kumar, S. 2020. Pre-training Tasks for Embedding-based Large-scale Retrieval. In *International Conference on Learning Representations*.
- Gutmann, M.; and Hyvärinen, A. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, 297–304.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9729–9738.
- Hoffmann, D. T.; Behrmann, N.; Gall, J.; Brox, T.; and Noroozi, M. 2022. Ranking Info Noise Contrastive Estimation: Boosting Contrastive Learning via Ranked Positives. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 897–905.
- Humeau, S.; Shuster, K.; Lachaux, M.-A.; and Weston, J. 2020. Poly-encoders: Architectures and Pre-training Strategies for Fast and Accurate Multi-sentence Scoring. In *International Conference on Learning Representations*.
- Jiang, X.; Liang, Y.; Chen, W.; and Duan, N. 2022. XLM-K: Improving Cross-Lingual Language Model Pre-training with Multilingual Knowledge. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 10840–10848.
- Johnson, J.; Douze, M.; and Jegou, H. 2021. Billion-Scale Similarity Search with GPUs. *IEEE Transactions on Big Data*, 535–547.
- Karpukhin, V.; Oguz, B.; Min, S.; Lewis, P. S.; Wu, L.; Edunov, S.; Chen, D.; and Yih, W.-t. 2020. Dense Passage Retrieval for Open-Domain Question Answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, 6769–6781.
- Khattab, O.; and Zaharia, M. 2020. ColBERT: Efficient and Effective Passage Search via Contextualized Late Interaction over BERT. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 39–48.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*.
- Li, S.; Lv, F.; Jin, T.; Lin, G.; Yang, K.; Zeng, X.; Wu, X.-M.; and Ma, Q. 2021. Embedding-based Product Retrieval in Taobao Search. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 3181–3189.
- Li, X.; Jin, J.; Zhou, Y.; Zhang, Y.; Zhang, P.; Zhu, Y.; and Dou, Z. 2025. From Matching to Generation: A Survey on Generative Information Retrieval. *ACM Transactions on Information Systems*, 1–62.
- Li, X.; Zhou, Y.; and Dou, Z. 2024. UniGen: A Unified Generative Framework for Retrieval and Question Answering with Large Language Models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 8688–8696.
- Lin, J.; Yadav, S.; Liu, F.; Rossi, N.; Suram, P. R.; Chembolu, S.; Chandran, P.; Mohapatra, H.; Lee, T.; Magnani, A.; et al. 2024. Enhancing Relevance of Embedding-based Retrieval at Walmart. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 4694–4701.
- Luan, Y.; Eisenstein, J.; Toutanova, K.; and Collins, M. 2021. Sparse, Dense, and Attentional Representations for Text Retrieval. *Transactions of the Association for Computational Linguistics*, 9: 329–345.
- Ma, X.; Wang, L.; Yang, N.; Wei, F.; and Lin, J. 2024. Fine-Tuning LLaMA for Multi-Stage Text Retrieval. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2421–2425.
- Magnani, A.; Liu, F.; Chaidaroon, S.; Yadav, S.; Reddy Suram, P.; Puthenpuhussery, A.; Chen, S.; Xie, M.; Kashi, A.; Lee, T.; et al. 2022. Semantic Retrieval at Walmart. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 3495–3503.
- Meghwani, H.; Agarwal, A.; Pattanayak, P.; Patel, H. L.; and Panda, S. 2025. Hard Negative Mining for Domain-Specific Retrieval in Enterprise Systems. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 6: Industry Track)*, 1013–1026.
- Oguz, B.; Lakhotia, K.; Gupta, A.; Lewis, P.; Karpukhin, V.; Piktus, A.; Chen, X.; Riedel, S.; Yih, S.; Gupta, S.; et al. 2022. Domain-matched Pre-training Tasks for Dense Retrieval. In *Findings of the Association for Computational Linguistics: NAACL 2022*, 1524–1534.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation Learning with Contrastive Predictive Coding. *arXiv preprint arXiv:1807.03748*.
- Qwen; Yang, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Li, C.; Liu, D.; Huang, F.; Wei, H.; Lin, H.; Yang, J.; Tu, J.; Zhang, J.; Yang, J.; Yang, J.; Zhou, J.; Lin, J.; Dang, K.; Lu, K.; Bao, K.; Yang, K.; Yu, L.; Li, M.; Xue, M.; Zhang, P.; Zhu, Q.; Men, R.; Lin, R.; Li, T.; Tang, T.; Xia, T.; Ren, X.; Ren, X.; Fan, Y.; Su, Y.; Zhang, Y.; Wan, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; and Qiu, Z. 2025. Qwen2.5 Technical Report. *arXiv preprint arXiv:2412.15115*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *International Conference on Machine Learning*, 8748–8763.
- Ren, R.; Lv, S.; Qu, Y.; Liu, J.; Zhao, W. X.; She, Q.; Wu, H.; Wang, H.; and Wen, J.-R. 2021. PAIR: Leveraging Passage-Centric Similarity Relation for Improving Dense Passage Retrieval. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 2173–2183.
- Robertson, S.; and Zaragoza, H. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. *Foundations and Trends® in Information Retrieval*, 333–389.

Rossi, N.; Lin, J.; Liu, F.; Yang, Z.; Lee, T.; Magnani, A.; and Liao, C. 2024. Relevance Filtering for Embedding-based Retrieval. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 4828–4835.

Su, W.; Ai, Q.; Li, X.; Chen, J.; Liu, Y.; Wu, X.; and Hou, S. 2024. Wikiformer: Pre-training with Structured Information of Wikipedia for Ad-hoc Retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 19026–19034.

Tang, Y.; Zhang, R.; Guo, J.; De Rijke, M.; Chen, W.; and Cheng, X. 2024. Listwise Generative Retrieval Models via a Sequential Learning Process. *ACM Transactions on Information Systems*, 42(5): 1–31.

Wang, Z.; Jiang, X.; Qin, Z.; and Yu, E. 2025. Personalized Query Auto-Completion for Long and Short-Term Interests with Adaptive Detoxification Generation. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, 5018–5028.

Wu, Z.; Xiong, Y.; Yu, S. X.; and Lin, D. 2018. Unsupervised Feature Learning via Non-Parametric Instance Discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3733–3742.

Xiong, L.; Xiong, C.; Li, Y.; Tang, K.-F.; Liu, J.; Bennett, P. N.; Ahmed, J.; and Overwijk, A. 2021. Approximate Nearest Neighbor Negative Contrastive Learning for Dense Text Retrieval. In *International Conference on Learning Representations*.

Yang, Z.; Shao, Z.; Dong, Y.; and Tang, J. 2024. TriSampler: A Better Negative Sampling Principle for Dense Retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 9269–9277.

Zhan, J.; Mao, J.; Liu, Y.; Guo, J.; Zhang, M.; and Ma, S. 2021. Optimizing Dense Retrieval Model Training with Hard Negatives. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1503–1512.

Zhang, S.; Liang, Y.; Gong, M.; Jiang, D.; and Duan, N. 2022. Multi-View Document Representation Learning for Open-Domain Dense Retrieval. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5990–6000.

Zheng, K.; Zhao, H.; Huang, R.; Zhang, B.; Mou, N.; Niu, Y.; Song, Y.; Wang, H.; and Gai, K. 2024. Full Stage Learning to Rank: A Unified Framework for Multi-Stage Systems. In *Proceedings of the ACM Web Conference 2024*, 3621–3631.

Zheng, Y.; Bian, J.; Meng, G.; Zhang, C.; Wang, H.; Zhang, Z.; Li, S.; Zhuang, T.; Liu, Q.; and Zeng, X. 2022. Multi-Objective Personalized Product Retrieval in Taobao Search. *arXiv preprint arXiv:2210.04170*.

Zhou, K.; Gong, Y.; Liu, X.; Zhao, W. X.; Shen, Y.; Dong, A.; Lu, J.; Majumder, R.; Wen, J.-R.; and Duan, N. 2022. SimANS: Simple Ambiguous Negatives Sampling for Dense Text Retrieval. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: Industry Track*, 548–559.