

Uni-Retrieval: A Multi-Style Retrieval Framework for STEM’s Education

Yanhao Jia ^{1*}, Xinyi Wu ^{2*}, Hao Li ³, Qinglin Zhang ²,
Yuxiao Hu ⁴, Shuai Zhao ^{1†}, Wenqi Fan ⁴

¹ Nanyang Technological University, Singapore;

² Shanghai Jiao Tong University, China; ³ Peking University, China;

⁴ Hong Kong Polytechnic University, Hong Kong, China.

shuai.zhao@ntu.edu.sg

Abstract

In AI-facilitated teaching, leveraging various query styles to interpret abstract text descriptions is crucial for ensuring high-quality teaching. However, current retrieval models primarily focus on natural text-image retrieval, making them insufficiently tailored to educational scenarios due to the ambiguities in the retrieval process. In this paper, we propose a diverse expression retrieval task tailored to educational scenarios, supporting retrieval based on multiple query styles and expressions. We introduce the **STEM Education Retrieval Dataset (SER)**, which contains over 24,000 query pairs of different styles, and the **Uni-Retrieval**, an efficient and style-diversified retrieval vision-language model based on prompt tuning. Uni-Retrieval extracts query style features as prototypes and builds a continuously updated Prompt Bank containing prompt tokens for diverse queries. This bank can be updated during test time to represent domain-specific knowledge for different subject retrieval scenarios. Our framework demonstrates scalability and robustness by dynamically retrieving prompt tokens based on prototype similarity, effectively facilitating learning for unknown queries. Experimental results indicate that Uni-Retrieval outperforms existing retrieval models in most retrieval tasks. This advancement provides a scalable and precise solution for diverse educational needs.

1 Introduction

Artificial Intelligence for Education (AI4EDU) is an emerging interdisciplinary field that leverages AI techniques to transform and enhance instructional design, learning processes, and assessment for education (Hwang et al., 2020). With the growing global emphasis on the importance of Science, Technology, Engineering, and Mathematics (STEM) education, retrieving

accurate resources from massive interdisciplinary knowledge databases has become a critical challenge.

Traditional retrieval systems are typically designed for natural text-image contents, which limits their utility in multi-modal STEM education contexts (Li et al., 2024a; Wang et al., 2023). Research indicates that these systems often fail to capture the complexity of materials such as images, diagrams, or interactive content, which are vital in STEM disciplines (Shen et al., 2023). Effective retrieval in STEM education should be able to handle different representations (i.e., different styles of text, image, audio, etc.) to accommodate the diverse learning and teaching needs within STEM subjects.

Despite advancements in text-image matching techniques (Williams-Lekuona et al., 2022; Zhou et al., 2024a), current retrieval systems still encounter challenges when implemented in STEM education (Li et al., 2025). These models are primarily optimized for matching text and images, neglecting the variety of query types essential in educational scenarios, including voice, sketches, and low-resolution images (Yang et al., 2023). The constraints of current frameworks within educational contexts frequently result in imprecise, biased, or inefficient retrieval outcomes, such deficiencies can impede teachers’ access to suitable instructional resources (Gasmi et al., 2024).

To address above challenges, we propose a multi-style and multi-modal retrieval task tailored for STEM education scenarios, as illustrated in Fig.1. The input types for this task include text, audio, and various styles of images, designed to meet the diverse needs of educational contexts. To adapt to this task, we introduce the **STEM Education Multi-Style Retrieval Dataset (SER)**, curated by 20 graduate and doctoral students from disciplines such as computer science, physics, energy, engineering,

*Equal Contribution, † Corresponding Author.

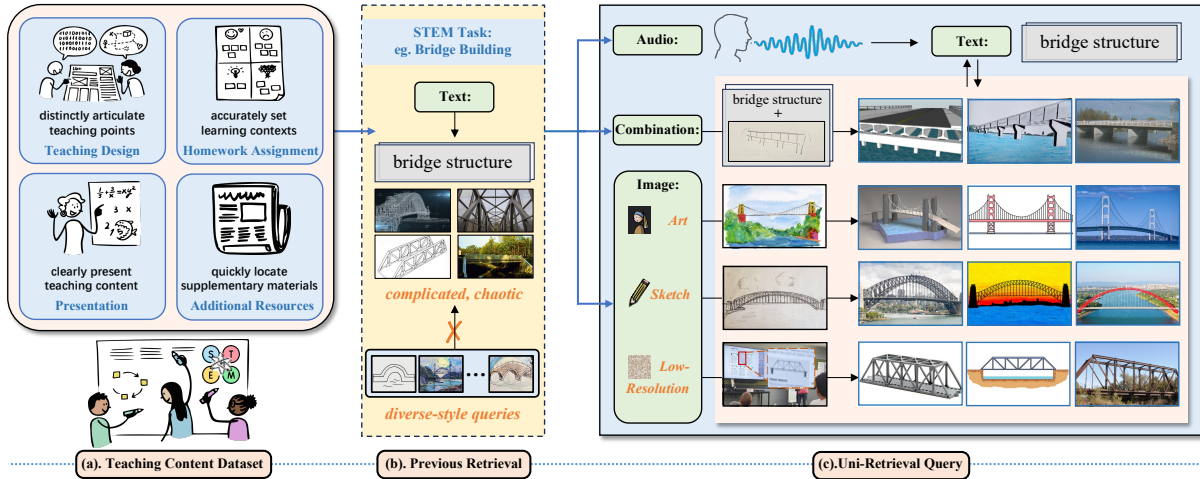


Figure 1: This advancement provides a scalable and precise solution for diverse educational needs. (b). Previous retrieval models focus on text-query retrieval data. (c) Our style-diversified retrieval setting accommodates the various query styles preferred by real educational content.

and mathematics. The dataset comprises 6,000 natural images with queries in various styles, including sketches, art, low-resolution images, text, and corresponding audio from different STEM fields. For modeling, we propose a novel plug-and-play feature representation structure called Prompt Bank, which provides a more universal representation of information across different disciplines and semantics by matching the abstract features of data. Building on Prompt Bank, we develop a lightweight retrieval model named **Uni-Retrieval**. Uni-Retrieval integrates Prompt Bank with various basic retrieval models, enabling efficient fine-tuning for educational scenarios. With only a small increase in parameters and inference time, Uni-Retrieval delivers significant performance improvements and fine-tuning capabilities, making it an ideal solution for STEM education retrieval tasks. Furthermore, it demonstrates the ability to perform effective retrieval across multiple unknown datasets, showcasing its scalability and adaptability in diverse scenarios. The main contributions of this work can be summarized as follows:

- To bridge the content gap between teacher and student or some abstract expression in STEM education field, we propose the multi-style multi-modal retrieval task for the STEM education. To adapt to this task, we construct the STEM Education Retrieval Dataset (SER), which contains several different subjects in STEM education.
- To efficiently and effectively train a model tailored to our task and dataset, we devise

the novel Uni-Retrieval algorithm, which incorporates a sustainably updatable Prompt Bank. Leveraging this bank, the Uni-Retrieval can represent different subjects at a high level and express the features of any other subjects by combining several prompts.

- Our Uni-Retrieval shows a more strong performance than any other previous method, not only on our SER dataset, but also on traditional retrieval dataset. We believe Uni-Retrieval can bring an infinite potential to STEM education community.

2 Preliminary

2.1 Task Formulation

We provide a formal problem formulation for query-based retrieval. Specifically, given an image I_i or a text prompt P_t from the style-specific query set Q_s , the retrieval model needs to compute the score between input and target queries and rank the corresponding answers A as high as possible. In the task settings for STEM education retrieval, which share a similar goal, the objective is to rank all answers correctly in response to input queries across various style-specific query sets $Q_{s=1}^n$. If the dataset does not contain the corresponding different style queries, the model should list the same category queries as the suggestions.

2.2 Dataset Construction

SER is a multi-style benchmark dataset we construct to facilitate accurate retrieval for teachers in STEM education. It contains a total of 24,000 text captions, audio clips and different

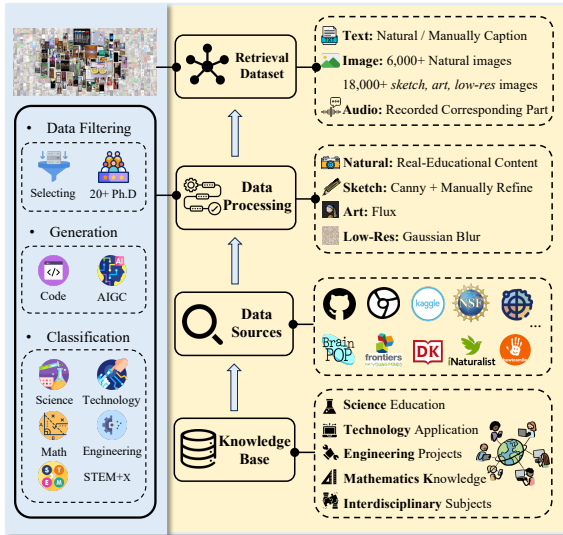


Figure 2: **Data construction pipeline.** 1. STEM education knowledge base. 2. Data sources: from online resources and dataset researches. 3. Data processing: extracting essential information from collected data, using AIGC algorithms to generate diverse modalities. 4. Retrieval dataset: construct total 24,000 images and multi-modal STEM educational dataset.

style queries to accommodate different educational scenarios. As illustrated in Fig.5, SER contains: **Text and Natural Image:** the most common query type, allowing teachers to describe problems using natural language or images for retrieval. **Audio:** a communication medium in education, enabling teachers to articulate complex queries, which can be further enhanced with LLMs or audio encoders. **Sketch:** hand-drawn sketches, whether created by users or written on blackboards, provide structural cues such as shape, pose, line, and edges to describe the problem. **Art:** art-style images as queries help bridge the gap between stylistically different images and original images, improving retrieval consistency across styles. **Low-Resolution:** queries involving lower-resolution images, such as those captured from a distance, ensure usability in scenarios where high-quality images are unavailable.

The details of the dataset construction pipeline are shown in Fig.2, we use the original STEM education image in the source dataset, and extensively collected datasets from the following sources: 1. online resources such as Kaggle, GitHub, BrainPOP, Frontiers, DKlearning, iNaturalist, AAES, etc. 2. relevant education dataset research, such as GAN (Jin et al., 2023) and PromptAloud (Lee et al., 2024b). To ensure high-quality data, more than 20 Ph.D. students from disciplines such as mathematics,

physics, chemistry, biology, electronics, computer science, and education conducted a secondary screening of the raw images. They also generate multi-modal combinations (image/text/audio) by leveraging AIGC models. Based on the natural images, the following steps were undertaken: **Text Generation:** we manually proofread the natural text descriptions. **Audio Recording:** the corresponding audio parts are recorded to match the text captions. **Sketch Images:** Canny algorithms are used to produce sketch images, Pidinet (Su et al., 2023) is employed to optimize and enhance low-quality sketches, and manual refinement is performed to achieve the final results. **Art-Style Images:** Flux model (Labs, 2023) is utilized to create art-style images. **Low-Resolution Images:** Gaussian Blur algorithms are applied to generate low-resolution images. According to the National Science Foundation (NSF)’s classification of STEM education, we collect 6,000 original samples spanning six styles, three modalities, and over 22 subjects. This comprehensive dataset, as illustrated in Appendix Fig.6, ensures diversity and quality across multiple educational domains.

3 Uni-Retrieval model

Our model consists of three main submodules: a **prototype learning module** for generating the prototype feature for each content style (Sec. 3.1), a **prompt bank** for saving features in the embedding space and selecting the prompt tokens which can represent the input’s style (Sec. 3.2), and a **feature extractor** based on different vision-language models for retrieving (Sec. 3.3). Additionally, we further present the training and inference process in Sec. 3.4.

3.1 Prototype Learning Module

For Uni-Retrieval, given an input query (freely based on audio, image, or text) $x \subseteq \mathbb{R}^{L \times C}$ and a feature encoder f , map x into a d -dimensional shared latent space (Prototype Module) using f , each style has m images. For style extracting, researchers usually use the pre-trained models which contains rich semantic information as the feature encoders. For example, if the input queries focus on style images, we leverage the style encoder (Tao, 2022) to extract visual features. If the query emphasizes the need for more textual information from the context, the text encoder and tokenizer, which are pretrained on large text datasets such as the Pile (Gao et al., 2020b), can be

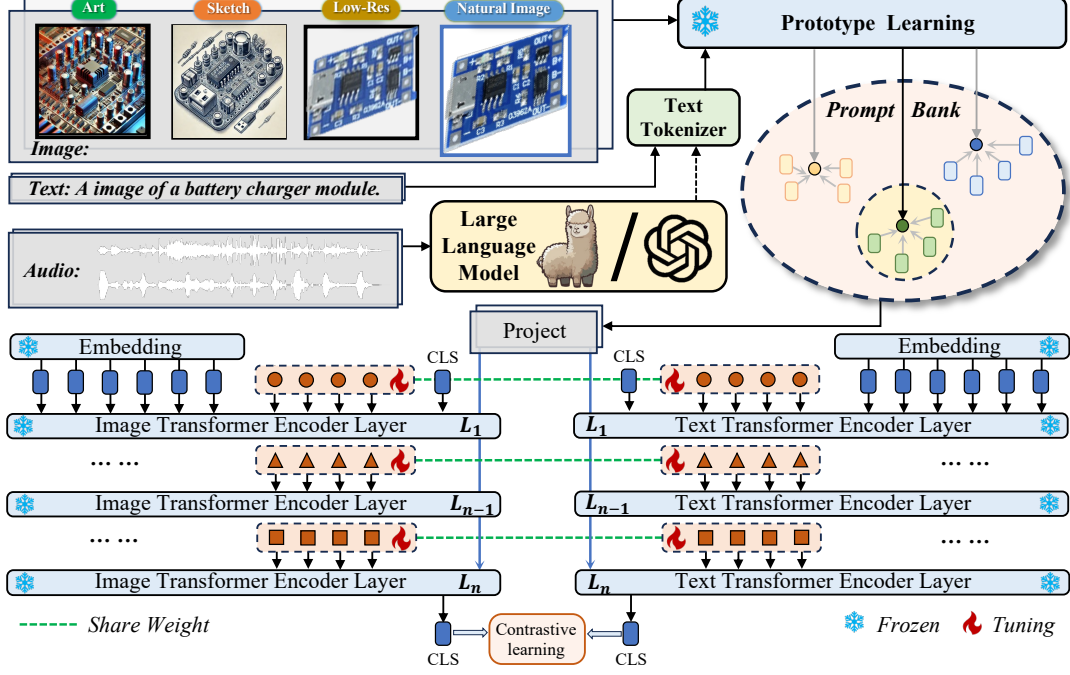


Figure 3: The Uni-Retrieval model's architecture.

utilized. The input query is embedded as follows:

$$E_0^i = f(x_0^i), \quad E_0^i \subseteq \mathbb{R}^d, i = 1, 2, \dots, m \quad (1)$$

where E_0^i denotes the 0-th style's i -th embedding. And then, using the average pooling operation to sum the different each style's embeddings to get the j -th prompt P_j :

$$P_j = \text{AvgPool}(\sum_{i=0}^m E_j^i) \quad (2)$$

3.2 Prompt Bank

The Prompt Bank builds a hidden state like TTT (Sun et al., 2024) and mamba (Gu and Dao, 2023), which are designed to store semantic and contextual information at a high level. Unlike the previous method, which leverages clusters to represent different styles, the Prompt Bank uses hash-like sets to store universal information. Prompt Bank contains N prompts, and each prompt k_i is associated as a value to a learnable key P_i :

$$\text{Prompt_Bank} = \{(k_1, P_1), \dots, (k_N, P_N)\} \quad (3)$$

We denote the key set as $K = \{k_i\}_{i=1}^N$ and define γ to score the match between k_i and P_i . Given an input x , the γ looks up the top- n keys to expand the feature embedding of x . The aim to use the hash-liked structure is promoting the matching speed between the input and the Prompt Bank's tokens. For Uni-Retrieval model, we calculate the cosine similarity between the matching prompt P_{j_i}

and the key K_i :

$$K_x = \arg \min_{\{j_i\}_{i=1}^n \subseteq [1, N]} \sum_{i=1}^n \gamma(P_{j_i}, K_i) \quad (4)$$

The Prompt Bank is free to combine different prompt tokens and expand feature embedding space, allowing different tokens associated with specific styles to jointly represent an input query. Due to the generalization properties on out-of-distribution of our Prompt Bank, even unseen styles also can combine similar styles' tokens to enhance retrieval performance by expand the semantic/context information provided by Prompt Bank. The expanding method is suitable for both different styles of images and different expression of text. Usually the special token is put at the beginning of the sequence tokens:

$$x_p = [CLS; P_{j_1}; P_{j_2}; \dots; P_{j_n}; x_e] \quad (5)$$

where x_p denotes the image's input feature after expanding prompt tokens; x_e represents the original sequence tokens patched from the input; CLS is the special token used for performing downstream tasks with different heads.

3.3 Feature Extractor

In Uni-Retrieval, we apply the ViT structure as the visual feature extractor, leverage a tokenizer to embed the input text query $x \subseteq \mathbb{R}^{L \times C}$, and then utilize the transformer (Vaswani et al., 2017) as the text encoder to extract features. The vision encoder

and the text encoder are initialized with OpenCLIP, and gpt-neo (Black et al., 2021) trained on the Pile dataset as the tokenizer, respectively. What’s more, we use GPT-4o (Hurst et al., 2024) as the optional choice of the external large language model to convert the audio clips to the text sequence.

The whole sequence tokens are fed into the feature extractor for training and inference layer by layer. Obtained from the Prompt Bank, visual prompt tokens represent various style information specific to different STEM subjects, while text prompt tokens convey distinct context information about different STEM subjects or different expressions about the same subjects. The parameters are shared between visual prompts and text prompts each layer to align vision and text modality. In the model training phase, the vision encoder, the tokenizer and the text encoder are fully frozen to retain the original semantic space.

3.4 Training and Inference

For the training procedure, in every training step, the style/context features are extracted from the corresponding encoder f to get the prompt P_j . Then, matching P_j and the key K_j to get the matching prompts P_j . Besides, the tokenizer and the patch layer map the inputs into sequence x_t :

$$x_t = \text{Tokenizer/Patch}(x) \quad (6)$$

where x_t denotes the temp state of features. After selecting n prompts following the aforementioned query strategy, the expanded feature embedding x_p is fed into the foundation model δ and getting the final result x_f . We use the CLS token to represent the whole sequence x_p following the settings of LLaMA (Touvron et al., 2023):

$$x_f = \delta(x_p)[:, 0, :] \quad (7)$$

The triplet loss function \mathcal{L} utilizes the features x_f , x_r , and x_h of an image or text, a retrieval target query, and a negative sample from a different category. Minimizing \mathcal{L} brings the correct sample pairs closer together while distancing the negative sample pairs. With μ as margin and distance function $d(a, b) = (1 - a * b) / (||a|| - ||b||)$, \mathcal{L} is given as:

$$\mathcal{L} = \max\{0, \mu + d(\delta(x_f), \delta(x_r)) - d(\delta(x_f), \delta(x_h))\} \quad (8)$$

where x_r, x_h denotes the embedding of the retrieval object and the negative sample respectively. Moreover, the key in Prompt Bank will be updated

with a scale parameter λ to weight the second loss function term:

$$\min_{k,p,L} \mathcal{L}(x_f, x_r, x_h) + \lambda \sum_{K_x} \gamma(q(x), k_{si}) \quad (9)$$

During the training procedure, Uni-Retrieval will jointly train both the Prompt Bank’s keys K and the prompt tokens P , which endow the Prompt Bank with the capability for continuous updates. The updatable parameters of the Uni-Retrieval model are limited compared to full-parameter fine-tuning, effectively saving computational resources and enhancing training speed.

For the inference procedure, we first extract the prototype feature for the unknown-style query input x . We use the prototype feature as the query to retrieve the fittest prompt tokens from the prompt bank. If the style is unknown type, the Prompt Bank will use several different clusters to represent it jointly. Then we prepend the prompt tokens x_p to the feature extractor δ for retrieval. The query embeddings are extracted in advance and stored in the databases to accelerate the retrieval process.

4 Experiments

4.1 Experiments Settings

We use our SER as the main experiment analysis and another three retrieval datasets to comprehensively evaluate the Uni-Retrieval’s performance. For evaluation metric, We evaluate the R@1 and R@5 metrics and the inference speed (ms) on all retrieval datasets. For R@1 and R@5, “ \uparrow ” denotes that higher is better. For ms, “ \downarrow ” denotes that quicker is better. Implement Details are detailed in Appendix C.

4.2 Comparison Experiment

On the SER dataset, Uni-Retrieval demonstrates superior performance across multiple scenarios with different query styles compared to other baselines, including multi-modality models, cross-modality pre-trained models, and prompt learning models. As shown in Tab.1 and Tab.4, the $T + S \rightarrow I$ means inputting the text and the style images as the multi-queries and outputting the corresponding images as the target queries. The experiment results yield three key observations:

The Uni-Retrieval achieves the best retrieval performance on the multi-style STEM Education Retrieval task: This highlights the effectiveness of Uni-Retrieval in handling complex multi-modal queries. Due to the Prompt

#	Method	Text → Image		Sketch → Image		Art → Image		Low-Res → Image	
		R@1↑	R@5↑	R@1↑	R@5↑	R@1↑	R@5↑	R@1↑	R@5↑
<i>Pretrained Cross-Modality Models</i>									
1	CLIP	54.6	78.4	47.3	68.9	46.8	71.3	53.7	72.9
2	BLIP	55.8	79.2	48.2	69.2	47.5	74.4	51.5	74.2
3	CLIP-Finetune	71.4	91.4	71.0	87.0	52.2	81.6	71.2	88.1
4	BLIP-Finetune	70.2	92.0	71.6	89.2	54.3	82.3	69.7	86.8
<i>Large Multi-Modality Models</i>									
5	LanguageBind	60.2	86.9	52.8	78.4	49.0	78.4	59.1	80.2
6	Unified-IO2	67.5	89.2	59.6	84.1	55.9	82.9	64.3	84.0
<i>Style Retrieval Models</i>									
7	SceneTrilogy	69.7	84.5	75.6	96.5	71.5	92.9	68.6	85.5
8	FashionNTM	50.4	81.3	68.9	88.6	67.1	88.9	45.6	77.5
<i>Cross-Modality Prompt Learning Models</i>									
9	VPT	69.9	84.1	53.3	72.3	62.7	83.2	67.4	79.1
10	CoCoOP	72.2	86.7	53.8	74.8	66.4	87.4	70.8	81.6
11	MaPLe	73.8	87.8	62.7	78.9	67.8	89.4	71.9	86.3
12	FreestyleRet	80.1	92.5	75.3	91.5	73.0	98.3	78.0	90.7
<i>Database-Driven Retrieval Models</i>									
13	GASKN	55.7	80.8	47.6	68.7	48.5	75.9	53.6	70.5
14	SKG	57.8	82.1	45.4	65.3	49.2	76.1	56.8	75.4
15	Uni-Retrieval	83.2	98.7	84.5	95.6	76.9	97.5	87.4	98.1

Table 1: Retrieval performance for STEM Education Retrieval task.

Method	Params(M)	Q2I(ms)↓	Q2T(ms)↓	T→I(Acc)↑
CLIP	427M	68ms	63ms	54.6
BLIP	891M	62ms	58ms	55.8
VPT	428M	73ms	69ms	69.9
LanguageBind	1200M	372ms	367ms	60.2
GASKN	33M	12ms	10ms	55.7
Uni-Retrieval	453M ₍₊₂₆₎	77ms ₍₊₉₎	70ms ₍₊₇₎	83.2 _(+28.6)

Table 2: The models inference speed comparison.

Method	T→I	T+S→I	I→T	I+S→T
CLIP-Finetune	54.6	55.3 _(+0.7)	47.4	46.6 _(-0.8)
VPT	69.9	72.0 _(+2.1)	73.9	74.1 _(+0.2)
Uni-Retrieval	83.2	87.4 _(+4.2)	81.7	83.3 _(+1.6)

Table 3: Retrieval performance with multi-style queries simultaneously on SER dataset.

Bank’s structure, Uni-Retrieval is a plug-and-play framework that can be highly flexible applied to various multi-modal models and enhance their retrieval capabilities. Line 15 in Tab.1 provides a substantial performance boost compared to both CLIP and CLIP-Finetune, further validating the effectiveness of our framework.

The Prototype module and Prompt Bank significantly outperform full-parameter fine-tuning: As shown in lines 3-4 and line 15 of Tab.1, Uni-Retrieval surpasses its fine-tuned CLIP counterpart by a large margin. Leveraging the prior knowledge bias introduced by the Prototype module and the efficient memory space of the Prompt Bank, Uni-Retrieval achieves superior results while tuning less than 5% of the model’s total parameters. This demonstrates the effectiveness of Uni-Retrieval’s design

in achieving high performance with minimal parameter adjustments.

Uni-Retrieval can simultaneously perform and mutually enhance traditional text-image retrieval performance: As shown in Tab.3, when handling text-image retrieval tasks, Uni-Retrieval allows multi-query inputs as additional references, significantly improving retrieval capability. This multi-query input design is not exclusive to Uni-Retrieval and can also benefit other retrieval models, offering a generalizable approach to enhancing retrieval tasks.

In addition to accuracy, inference speed is a crucial metric for evaluating retrieval models. As shown in Tab.2, Uni-Retrieval adds just 9ms per search iteration. Compared to GASKN, Uni-Retrieval demonstrates significantly stronger retrieval performance than traditional database-driven methods. Additionally, when compared to other cross-modality methods, Uni-Retrieval excels in both tuning efficiency and retrieval accuracy, further validating its effectiveness and scalability.

4.3 Ablation Study

To quantitatively evaluate the role of prompts in the model, we conducted ablation studies on the insertion type and token number of prompt tokens within Uni-Retrieval for four style metrics. These studies aimed to assess their impact on the style-diversified STEM education retrieval task,

#	Method	I→T	S→T	A→T	L→T	T→S	T→A	T→L	S→A	S→L	A→L
<i>Metric: R@I↑ on SER Dataset</i>											
1	CLIP	47.4	38.4	37.9	38.6	38.8	37.4	35.7	36.9	34.8	31.5
2	BLIP	48.9	39.2	38.4	39.5	39.7	37.1	36.5	35.0	34.9	32.6
3	CLIP-Finetune	75.7	72.4	71.3	69.8	70.2	68.5	67.7	65.4	66.8	66.3
4	BLIP-Finetune	77.4	73.0	72.6	70.5	71.3	69.4	68.1	66.2	67.2	67.0
5	LanguageBind	55.4	54.9	53.1	53.4	49.7	48.7	49.1	46.2	46.8	45.9
6	Unified-IO2	57.3	57.2	56.3	54.5	51.1	49.9	48.6	48.0	47.2	46.8
7	SceneTrilogy	72.4	76.5	70.6	71.5	69.3	69.9	68.7	65.2	66.2	64.4
8	FashionNTM	70.6	73.3	68.9	69.6	67.1	68.0	66.5	67.5	64.8	62.4
9	VPT	73.9	71.8	70.4	68.7	69.0	68.2	67.4	66.6	64.5	63.8
10	CoCoOP	76.5	74.7	73.4	74.0	71.4	72.3	70.8	68.9	67.2	67.3
11	MaPLe	78.3	75.8	75.7	74.9	72.4	69.6	69.2	68.3	67.4	65.6
12	FreestyleRet	80.8	73.5	75.5	71.4	73.0	68.3	68.0	69.4	70.6	68.9
13	GASKN	53.8	52.9	52.6	50.7	49.4	47.9	46.0	47.1	47.3	45.9
14	SKG	54.3	51.7	50.4	51.3	48.5	46.1	45.4	46.9	47.0	45.9
15	Uni-Retrieval	81.7	76.3	74.9	77.6	73.5	74.2	78.0	71.4	72.3	70.8

Table 4: Retrieval performance for STEM Education Retrieval task.

#	Type	Token-Num	T→I	S→I	A→I	L→I
1	Deep	1	72.0	78.3	73.2	80.6
2	Deep	2	77.1	81.2	75.5	85.8
3	Deep	8	83.24	82.7	76.5	87.0
4	Shallow	4	68.2	75.6	70.4	77.3
5	Deep	4	83.2	84.5	76.9	87.4

Table 5: Ablation study for prompt settings.

providing insights into how prompts influence performance and model adaptability. The shallow type involves inserting prompt tokens only at the first layer, while the deep type inserts prompt tokens across all layers. The token number refers to the number of repeated prompts.

As shown in lines 4-5 in Tab.5, the results indicate that the deep type consistently outperforms the shallow type. Additionally, lines 1-3 and line 5 in Tab.5 demonstrate the change on number of prompt tokens. We observed that repeating the prompt tokens more than four times does not significantly improve model performance. Instead, it rapidly increases the number of tuning parameters, which slows down both the tuning process and inference speed. This indicates that four repetitions provide a balanced trade-off between performance and computational efficiency. Therefore, we ultimately selected four prompt tokens to be inserted at each layer as the standard configuration for Uni-Retrieval. This choice effectively balances model performance, tuning efficiency, and inference speed, which can serve as a valuable reference for other tuning models.

We also evaluated Uni-Retrieval’s zero-shot retrieval performance on several other multi-style datasets. As shown in Tab.6, we compared Uni-Retrieval against various baseline models across three datasets: the DSR, DomainNet,

Method	Text → Im.		Sketch → Im.		Art → Im.	
	R@1↑	R@5↑	R@1↑	R@5↑	R@1↑	R@5↑
<i>Diverse-Style Retrieval Dataset</i>						
LanguageBind	71.0	95.5	50.8	79.4	58.2	86.3
CoCoOP	71.4	94.6	77.5	97.2	69.3	97.1
MaPLe	73.1	95.9	80.3	97.9	70.6	97.2
FreestyleRet	71.4	97.2	81.6	98.0	72.3	98.1
Uniretrieval	82.3	97.4	82.7	98.9	75.1	98.0
<i>DomainNet Dataset</i>						
VPT	59.7	86.1	53.5	77.3	54.6	81.8
BLIP-Finetune	65.3	94.2	71.4	89.7	54.3	82.3
FreestyleRet	70.2	95.2	75.2	93.2	73.1	92.6
Uniretrieval	70.7	96.0	77.6	94.1	73.4	92.9
<i>SketchCOCO Dataset</i>						
MaPLe	26.4	53.9	18.0	48.3	-	-
SceneTrilogy	30.6	65.8	22.5	51.6	-	-
FreestyleRet	31.5	67.3	29.6	56.1	-	-
Uniretrieval	34.7	71.6	30.2	60.4	-	-

Table 6: The zero-shot retrieval performance comparison on retrieval datasets.

and SketchCOCO dataset, each representing distinct domains of style-based queries. As shown in Tab.6, Uni-Retrieval demonstrates exceptional zero-shot performance across these diverse datasets, highlighting its capability to perform effective information retrieval in various previously unknown databases. This performance underscores Uni-Retrieval’s scalability and robustness, significantly enhancing the adaptability and effectiveness of existing retrieval models in handling diverse and unstructured data domains.

4.4 Visualization Result

In Fig.4, we visualize the style-diversified query inputs and their corresponding retrieval answers from our Uni-Retrieval and the FreestyleRet baseline model. We summarize three common retrieval errors in the case analysis, where

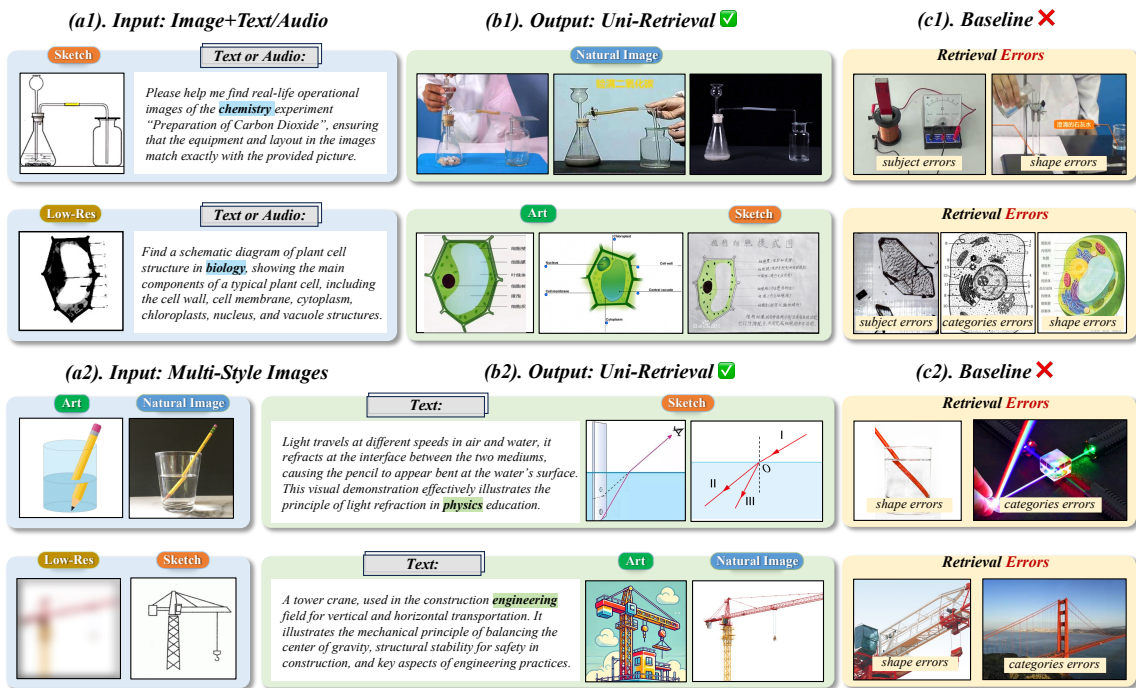


Figure 4: The case study for our Uni-Retrieval and the FreestyleRet baseline.

subject errors, semantic errors, and color errors represent the false retrieval result with false subjects, semantic information, and colors. We propose the subject error cases in Fig.4(a1)-(c1). The subject information is contained widely in different style queries. Thus, pose error cases occur in sketch, art, and low-resolution queries. Subject information is conveyed through the primary objects in images and their corresponding textual descriptions. Subject errors occur when there is incorrect recognition or classification of these objects, leading to mismatched associations between the image and text.

Semantic errors, on the other hand, arise from inaccuracies in describing object details. These errors frequently occur when irrelevant text is associated with specific parts of an object, particularly in the context of art descriptions. Such mismatches result in the model generating incorrect attention maps, thereby failing to accurately connect the visual and textual elements. Thus, in Fig.4(a2)-(c2), most of the semantic errors occur in the art-style retrieval task.

For the low-resolution query retrieval task, color is vital retrieval information. We show the color errors from the low-resolution retrieval task. Compared to the FreestyleRet baseline model, our Uni-Retrieval achieves fine-grained retrieval based on subject, semantic, and color information from style-diversified query inputs. It demonstrates a superior understanding of

semantic information and fine-grained alignment between modalities, particularly in the precise description and representation of key object parts. This highlights the significant advantages and capabilities of our Uni-Retrieval framework.

5 Conclusion

To address the challenge of fine-grained and efficient retrieval in STEM teaching scenarios, we proposed a multi-style and multi-modal STEM education retrieval task and curated a multi-style dataset of over 24,000 samples for model fine-tuning. To balance training efficiency and retrieval performance, we developed a lightweight and plug-and-play feature expression module, Prompt Bank, and built a database-driven accurate retrieval model, Uni-Retrieval, based on Prompt Bank. Compared to current state-of-the-art retrieval models, Uni-Retrieval significantly improves retrieval performance with only a 26M (less than 5%) increase in parameter size and less than 10ms additional retrieval time. Furthermore, the training and deployment costs of Uni-Retrieval are substantially lower than those of existing large retrieval models, making it a more economical and practical solution for educational scenarios. We hope Uni-Retrieval can inspire new possibilities for the community, offering an effective and accessible approach to retrieval in STEM education and beyond.

Limitation

However, our work still has some limitations that require further research. Firstly, STEM education differs significantly from higher education, K-12 education, humanities education, and other scenarios in terms of data and usage requirements. A key challenge for future research is how to maintain efficient retrieval performance while adapting to a wider range of educational scenarios. Additionally, we plan to exploring how to efficiently acquire various professional educational knowledge based on VLMs. These improvements aim to make Uni-Retrieval more versatile and impactful across diverse educational domains.

References

- Yehia Ibrahim Alzoubi, Ahmet Ercan Topcu, and Erdem Ozdemir. 2024. Enhancing document image retrieval in education: Leveraging ensemble-based document image retrieval systems for improved precision. *Applied Sciences*, 14(2):751.
- Eugene Belilovsky, Matthew Blaschko, Jamie Ryan Kiros, Raquel Urtasun, and Richard Zemel. 2017. Joint embeddings of scene graphs and images. In *International Conference On Learning Representations-Workshop*.
- Sid Black, Gao Leo, Phil Wang, Connor Leahy, and Stella Biderman. 2021. [GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow](#).
- T Brown, B Mann, N Ryder, M Subbiah, JD Kaplan, P Dhariwal, A Neelakantan, P Shyam, G Sastry, A Askell, et al. 2020. Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*.
- Wei Chen, Yu Liu, Weiping Wang, Erwin M Bakker, Theodoros Georgiou, Paul Fieguth, Li Liu, and Michael S Lew. 2022. Deep learning for instance retrieval: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(6):7270–7292.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, et al. 2024. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 24185–24198.
- Eulrang Cho, Jooyeon Kim, and Hyunwoo J Kim. 2023. Distribution-aware prompt tuning for vision-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22004–22013.
- Youngduck Choi, Younghan Lee, Dongmin Shin, Junghyun Cho, Seoyon Park, et al. 2020. Ednet: A large-scale hierarchical dataset in education. In *Artificial Intelligence in Education: 21st International Conference, AIED 2020*, pages 69–73.
- Pinaki Nath Chowdhury, Ayan Kumar Bhunia, Aneeshan Sain, Subhadeep Koley, Tao Xiang, and Yi-Zhe Song. 2023a. Scenetriology: On human scene-sketch and its complementarity with photo and text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10972–10983.
- Pinaki Nath Chowdhury, Ayan Kumar Bhunia, Aneeshan Sain, Subhadeep Koley, Tao Xiang, and Yi-Zhe Song. 2023b. Scenetriology: On human scene-sketch and its complementarity with photo and text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10972–10983.
- Pinaki Nath Chowdhury, Aneeshan Sain, Ayan Kumar Bhunia, Tao Xiang, Yulia Gryaditskaya, and Yi-Zhe Song. 2022. Fs-coco: Towards understanding of freehand sketches of common objects in context. In *European conference on computer vision*, pages 253–270. Springer.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, et al. 2024. A survey on in-context learning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1107–1128.
- Abolfazl Farahani, Sahar Voghoei, Khaled Rasheed, and Hamid R Arabnia. 2021. A brief review of domain adaptation. *Advances in data science and information engineering: proceedings from ICDATA 2020 and IKE 2020*, pages 877–894.
- Chengying Gao, Qi Liu, Qi Xu, Limin Wang, Jianzhuang Liu, and Changqing Zou. 2020a. Sketchycoco: Image generation from freehand scene sketches. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5174–5183.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, et al. 2020b. The Pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Karim Gasmi, Hajer Ayadi, and Mouna Torjmen. 2024. Enhancing medical image retrieval with umls-integrated cnn-based text indexing. *Diagnostics*, 14(11):1204.
- Chunjiang Ge, Rui Huang, Mixue Xie, Zihang Lai, Shiji Song, Shuang Li, and Gao Huang. 2023. Domain adaptation via prompt learning. *IEEE Transactions on Neural Networks and Learning Systems*.

- Albert Gu and Tri Dao. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, et al. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- Jinchang Hou, Chang Ao, Haihong Wu, Xiangtao Kong, et al. 2024. E-eval: A comprehensive chinese k-12 education evaluation benchmark for large language models. *arXiv preprint arXiv:2401.15927*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Gwo-Jen Hwang, Haoran Xie, Benjamin W Wah, and Dragan Gašević. 2020. Vision, challenges, roles and research issues of artificial intelligence in education.
- Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, et al. 2022. Visual prompt tuning. In *European Conference on Computer Vision*, pages 709–727.
- Yanhao Jia, Xinyi Wu, Qinglin Zhang, Yiran Qin, Luwei Xiao, and Shuai Zhao. 2025a. [Towards robust evaluation of stem education: Leveraging mllms in project-based learning](#). *ResearchGate*.
- Yanhao Jia, Ji Xie, S Jivaganesh, Hao Li, Xu Wu, and Mengmi Zhang. 2025b. Seeing sound, hearing sight: Uncovering modality bias and conflict of ai models in sound localization. *arXiv preprint arXiv:2505.11217*.
- Yuxi Jin, Ping Li, Wenxiao Wang, Suiyun Zhang, Di Lin, and Chengjiu Yin. 2023. Gan-based pencil drawing learning system for art education on large-scale image datasets with learning analytics. *Interactive Learning Environments*, 31(5):2544–2561.
- Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. 2015. Image retrieval using scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3668–3678.
- Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, et al. 2023. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19113–19122.
- Mert Kilickaya and Arnold WM Smeulders. 2021. Structured visual search via composition-aware learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1701–1710.
- Iasonas Kokkinos. 2017. Ubernet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6129–6138.
- Soonwoo Kwon, Sojung Kim, Minju Park, Seunghyun Lee, and Kyuseok Kim. 2024. Biped: Pedagogically informed tutoring system for esl education. *arXiv preprint arXiv:2406.03486*.
- Black Forest Labs. 2023. Flux. <https://github.com/black-forest-labs/flux>.
- Saehyung Lee, Sangwon Yu, Junsung Park, Jihun Yi, and Sungroh Yoon. 2024a. Interactive text-to-image retrieval with large language models: A plug-and-play approach. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, pages 791–809.
- Unggi Lee, Ariel Han, Jeongjin Lee, Eunseo Lee, Jiwon Kim, Hyeoncheol Kim, and Cheolil Lim. 2024b. Prompt aloud!: Incorporating image-generative ai into steam class with learning analytics using prompt data. *Education and Information Technologies*, 29(8):9575–9605.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021a. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021b. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059.
- Ang Li, Jin Sun, Joe Yue-Hei Ng, Ruichi Yu, Vlad I Morariu, and Larry S Davis. 2017. Generating holistic 3d scene abstractions for text-based image retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 193–201.
- Hao Li, Jinfa Huang, Peng Jin, Guoli Song, Qi Wu, and Jie Chen. 2023a. Weakly-supervised 3d spatial reasoning for text-based visual question answering. In *Transactions on Image Processing*, pages 3367–3382. IEEE.
- Hao Li, Yanhao Jia, Peng Jin, Zesen Cheng, Kehan Li, Jialu Sui, Chang Liu, and Li Yuan. 2025. Freestyleret: Retrieving images from style-diversified queries. In *European Conference on Computer Vision*, pages 258–274.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023b. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR.
- Tieying Li, Xiaochun Yang, Yiping Ke, Bin Wang, Yinan Liu, and Jiaying Xu. 2024a. Alleviating the inconsistency of multimodal data in cross-modal retrieval. In *2024 IEEE 40th International Conference on Data Engineering*, pages 4643–4656.

- Wen Li, Lixin Duan, Dong Xu, and Ivor Wai-Hung Tsang. 2011. Text-based image retrieval using progressive multi-instance learning. In *2011 international conference on computer vision*, pages 2049–2055. IEEE.
- Yongqi Li, Wenjie Wang, Leigang Qu, Liqiang Nie, Wenjie Li, and Tat-Seng Chua. 2024b. Generative cross-modal retrieval: Memorizing images in multimodal language models for retrieval and beyond. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, pages 11851–11861.
- Do Long, Yiran Zhao, Hannah Brown, Yuxi Xie, James Zhao, Nancy Chen, Kenji Kawaguchi, Michael Shieh, and Junxian He. 2024. Prompt optimization via adversarial in-context learning. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, pages 7308–7327.
- Jiasen Lu, Christopher Clark, Sangho Lee, Zichen Zhang, Savya Khosla, Ryan Marten, Derek Hoiem, and Aniruddha Kembhavi. 2024. Unified-io 2: Scaling autoregressive multimodal models with vision language audio and action. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26439–26455.
- Long Mai, Hailin Jin, Zhe Lin, Chen Fang, Jonathan Brandt, and Feng Liu. 2017. Spatial-semantic image search by visual feature synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4718–4727.
- Andrei Neculai, Yanbei Chen, and Zeynep Akata. 2022. Probabilistic compositional embeddings for multimodal image retrieval. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4547–4557.
- Xing Nie, Bolin Ni, Jianlong Chang, Gaofeng Meng, Chunlei Huo, et al. 2023. Pro-tuning: Unified prompt tuning for vision tasks. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Ankit Pal, Logesh Kumar Umaphathi, and Malaikannan Sankarasubbu. 2022. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. In *Conference on health, inference, and learning*, pages 248–260. PMLR.
- Anwesan Pal, Sahil Wadhwa, Ayush Jaiswal, Xu Zhang, Yue Wu, Rakesh Chada, Pradeep Natarajan, et al. 2023. Fashionntm: Multi-turn fashion image retrieval via cascaded memory. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11323–11334.
- Harry A Patrinos and Noam Angrist. 2018. Global dataset on education quality: A review and update (2000-2017). *World Bank Policy Research Working Paper*, (8592).
- Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. 2019. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1406–1415.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763.
- S Ruder. 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*.
- Sheng Shen, Shijia Yang, Tianjun Zhang, Bohan Zhai, Joseph E Gonzalez, Kurt Keutzer, and Trevor Darrell. 2024. Multitask vision-language prompt tuning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5656–5667.
- Tao Shen, Xiubo Geng, Chongyang Tao, Can Xu, Guodong Long, Kai Zhang, and Daxin Jiang. 2023. Unifier: A unified retriever for large-scale retrieval. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4787–4799.
- Karen Simonyan. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Zhuo Su, Jiehua Zhang, Longguang Wang, Hua Zhang, Zhen Liu, Matti Pietikäinen, and Li Liu. 2023. Lightweight pixel difference networks for efficient visual representation learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(12):14956–14974.
- Yu Sun, Xinhao Li, Karan Dalal, Jiarui Xu, Arjun Vikram, Genghan Zhang, Yann Dubois, Xinlei Chen, Xiaolong Wang, et al. 2024. Learning to (learn at test time): Rnns with expressive hidden states. *Preprint*, arXiv:2407.04620.
- Yilin Tao. 2022. Image style transfer based on vgg neural network model. In *2022 IEEE International Conference on Advances in Electrical Engineering and Computer Applications (AEECA)*, pages 1475–1482.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30.
- Jinzheng Wang, Xin Liang, Ben Whitney, Jieyang Chen, Qian Gong, Xubin He, et al. 2023. Improving progressive retrieval for hpc scientific data using deep neural network. In *2023 IEEE 39th International Conference on Data Engineering*, pages 2727–2739.

- Shen Wang, Tianlong Xu, Hang Li, Chaoli Zhang, Joleen Liang, Jiliang Tang, Philip S Yu, and Qingsong Wen. 2024a. Large language models for education: A survey and outlook. *arXiv preprint arXiv:2403.18105*.
- Sijin Wang, Ruiping Wang, Ziwei Yao, Shiguang Shan, and Xilin Chen. 2020. Cross-modal scene graph matching for relationship-aware image-text retrieval. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1508–1517.
- Yidong Wang, Zhuohao Yu, Wenjin Yao, Zhengran Zeng, Linyi Yang, Cunxiang Wang, et al. 2024b. Pandalm: An automatic evaluation benchmark for llm instruction tuning optimization. In *The Twelfth International Conference on Learning Representations*.
- Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. 2022. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 139–149.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Johannes Welbl, Nelson F Liu, and Matt Gardner. 2017. Crowdsourcing multiple choice science questions. *arXiv preprint arXiv:1707.06209*.
- Mikel Williams-Lekuona, Georgina Cosma, and Iain Phillips. 2022. A framework for enabling unpaired multi-modal learning for deep cross-modal hashing retrieval. *Journal of Imaging*, 8(12):328.
- Luwei Xiao, Rui Mao, Shuai Zhao, Qika Lin, Yanhao Jia, Liang He, and Erik Cambria. 2025. Exploring cognitive and aesthetic causality for multimodal aspect-based sentiment analysis. *IEEE Transactions on Affective Computing*.
- Qu Yang, Mang Ye, Zhaohui Cai, Kehua Su, and Bo Du. 2023. Composed image retrieval via cross relation network with hierarchical aggregation transformer. *IEEE Transactions on Image Processing*.
- Yongxin Yang and Timothy M Hospedales. 2016. Trace norm regularised deep multi-task learning. *arXiv preprint arXiv:1606.04038*.
- Yang Yu, Meiyu Liang, Mengran Yin, Kangkang Lu, Junping Du, and Zhe Xue. 2024. [Unsupervised multimodal graph contrastive semantic anchor space dynamic knowledge distillation network for cross-media hash retrieval](#). In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, pages 4699–4708.
- Yu Zhang and Qiang Yang. 2022. [A survey on multi-task learning](#). *IEEE Transactions on Knowledge and Data Engineering*, 34(12):5586–5609.
- Shuai Zhao, Meihuizi Jia, Luu Anh Tuan, Fengjun Pan, and Jinming Wen. 2024a. Universal vulnerabilities in large language models: Backdoor attacks for in-context learning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11507–11522.
- Shuai Zhao, Jinming Wen, Anh Luu, Junbo Zhao, and Jie Fu. 2023. Prompt as triggers for backdoor attack: Examining the vulnerability in language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12303–12317.
- Shuai Zhao, Xiaobao Wu, Cong-Duy Nguyen, Meihuizi Jia, Yichao Feng, and Luu Anh Tuan. 2024b. [Unlearning backdoor attacks for llms with weak-to-strong knowledge distillation](#). *arXiv preprint arXiv:2410.14425*.
- Juncheng Zheng, Meiyu Liang, Yang Yu, Yawen Li, and Zhe Xue. 2024. [Knowledge graph enhanced multimodal transformer for image-text retrieval](#). In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, pages 70–82.
- Hongliang Zhou, Yufan Hu, Shuai Liu, Guoxiong Zhou, Jiabin Xu, Aibin Chen, Yanfeng Wang, Liujuan Li, and Yahui Hu. 2024a. A precise framework for rice leaf disease image-text retrieval using fhtw-net. *Plant Phenomics*, 6:0168.
- Junjie Zhou, Zheng Liu, Shitao Xiao, Bo Zhao, and Yongping Xiong. 2024b. VISTA: Visualized text embedding for universal multi-modal retrieval. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, pages 3185–3200.
- Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. 2022a. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4396–4415.
- Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. 2022b. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16816–16825.
- Bin Zhu, Bin Lin, Munan Ning, Yang Yan, Jiayi Cui, et al. 2023. Languagebind: Extending video-language pretraining to n-modality by language-based semantic alignment. *arXiv preprint arXiv:2310.01852*.

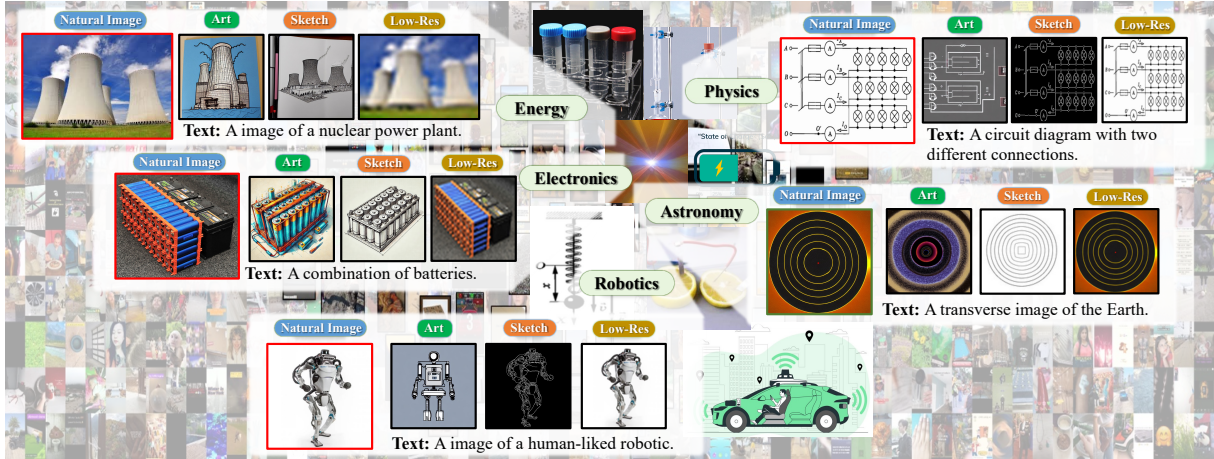


Figure 5: **The SER Dataset** contains 24,000+ text captions and their corresponding queries with various styles, including Natural, Sketch, Art, Low-Resolution (Low-Res) images and audio clips from different STEM subjects.

A Related Works

A.1 Dataset Adaptation in Education

Within the realm of education, image retrieval in education has distinct characteristics, as images often reflect the teaching intentions of educators. This facilitates the rapid and accurate alignment of visual content with teaching materials, thereby reducing educators’ preparation workload and enhancing the precision of learning data. While existing researches have focused on classifying educational data (Choi et al., 2020), they often encounter constraints. Due to the complexity of incorporating an expansive range of teaching scenarios in STEM education and the scarcity of data, numerous studies often narrow the scope to a limited set of subject applications (Hendrycks et al., 2021; Pal et al., 2022) or to a limited set of teaching strategy retrievals (Kwon et al., 2024; Welbl et al., 2017).

There is a considerable variation across existing STEM education datasets regarding its specific composition. Many datasets are cluttered with irrelevant or invalid data, lack comprehensive coverage of specialised content, and suffer from quality assurance issues (Patrinos and Angrist, 2018). Although STEM education datasets are assembled from interactions between learners and large language models (Hou et al., 2024; Wang et al., 2024a), they are generally not well-suited for use by educators and learners across multiple domains. Furthermore, creating a precise and professional data retrieval repository for the educational domain requires efficient retrieval algorithms as support (Alzoubi et al., 2024). To ensure efficient retrieval and usability in STEM

education scenarios, we construct the SER dataset, which includes multiple query styles to enhance retrieval diversity.

A.2 Multi-task Learning

In STEM education, the multi-style retrieval model needs to leverage multi-task learning to align features and learning across different modal samples. Multi-task learning refers to the simultaneous training and optimization of multiple related tasks within a single model (Zhang and Yang, 2022; Xiao et al., 2025; Jia et al., 2025b,a). By sharing parameters and representations across functions, it improves overall performance. Compared to other transfer learning methods, including domain adaptation (Farahani et al., 2021) and domain generalization (Zhou et al., 2022a), multi-task learning annotates data and achieves CLIP-level model fine-tuning and convergence, the data of each task in multi-task learning is well-labeled.

Overall, multi-task learning introduces a new tool for STEM education practitioners that may help meet requirements, especially if speed and efficiency are preferred over performance. While many recent multi-task learning employ two clusters of contemporary techniques, hard parameter sharing and soft parameter sharing (Ruder, 2017). In hard parameter sharing, most or all of the parameters in the network are shared among all tasks (Kokkinos, 2017). In soft parameter sharing, the models are tied together either by information sharing or by requiring parameters to be similar (Yang and Hospedales, 2016). Consequently, our Uni-Retrieval adopts a blended multi-task learning paradigm, adopt

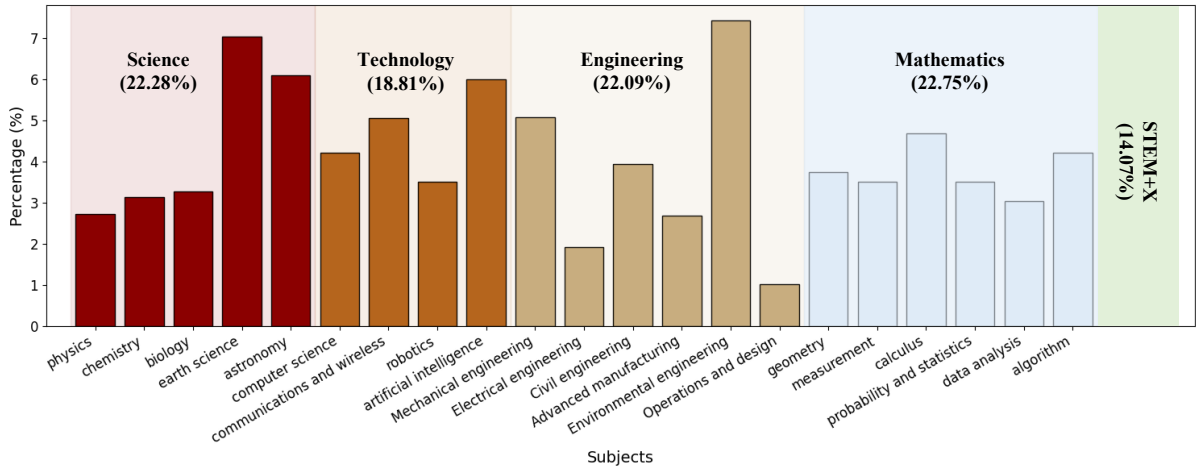


Figure 6: **Concept distribution of our SER dataset.** Our dataset exhibits a diverse distribution on different concept domains.

both hard and soft parameter in different styles of tasks. Building upon successful multi-task learning method for CLIP, such as CoCoOP (Zhou et al., 2022b), MaPLe (Khattak et al., 2023), and FreestyleRet (Li et al., 2025), our study leverages these techniques to strengthen domain adaptation and multi-task learning.

A.3 Query-based Retrieval

Existing work in Query-based Image Retrieval (QBIR) primarily includes content-based image retrieval (Chen et al., 2022), text-based image retrieval (Li et al., 2011), and multi-modal retrieval (Neculai et al., 2022). In content-based image retrieval, the visual features of images are directly utilized for retrieval. However, its reliance on fixed content and location makes it relatively inflexible in capturing diverse user intents (Lee et al., 2024a). Alternative methods like sketching (Chowdhury et al., 2022, 2023a) and scene graph construction (Johnson et al., 2015) enable the retrieval of abstract images that are hard to describe verbally, though they lack the intuitive ease of natural language-based retrieval. In text-based image retrieval, enhancements to text queries often involve indicating content structure. However, these approaches are either restricted by closed vocabularies (Mai et al., 2017; Kilickaya and Smeulders, 2021) or face substantial challenges (Li et al., 2017) in deriving structures from natural language descriptions. Recent multi-modal approaches, such as cross-modal scene graph-based image-text retrieval (Wang et al., 2020) and joint visual-scene graph embedding for image retrieval (Belilovsky et al., 2017), still depend on word

embeddings and image features.

Despite advancements in QBIR, challenges including the semantic gap that can lead to inaccurate retrieval results, high computational complexity and resource costs for large-scale image databases, and the high cost of obtaining quality data annotations (Li et al., 2024b). The application of QBIR to educational resource retrieval is promising but has been hindered by the complexity of educational discourse, the limitations of educational databases, and the associated costs (Zhou et al., 2024b). Our query model effectively combines multi-modal retrieval methods, integrating audio and natural language with multi-style image inputs. The former enables natural and rapid expression of content, while the latter facilitates accurate and intuitive image localization, enhancing educational data retrieval.

A.4 Prompt Tuning

Prompt tuning (Brown et al., 2020) was first proposed in natural language processing (NLP) and has been an efficient approach that bridges the gap between pre-trained language models and downstream tasks (Li et al., 2023a; Zhao et al., 2023, 2024b). Prompt tuning leverages natural language prompts to optimize the language model’s ability to understand tasks, which demonstrates exceptional performance in few-shot and zero-shot learning. Recent studies have focused on optimizing various components of prompt tuning, such as prompt generation, continuous prompt optimization (Lester et al., 2021a), and adapting to large-scale models through methods like in-context learning (Dong et al., 2024; Zhao et al., 2024a),

instruction-tuning (Wang et al., 2024b), and chain-of-thought (Wei et al., 2022). For example, Lester et al. (2021b) leverage soft prompts to condition frozen language models to enhance the performance of specific downstream tasks. Long et al. (2024) propose an adversarial in-context learning algorithm, which leverages adversarial learning to optimize task-related prompts.

Furthermore, prompt tuning has gradually become a pivotal technique in computer vision (Shen et al., 2024), enabling efficient adaptation of pre-trained models to diverse tasks. Notable methods include visual prompt tuning for classification (Jia et al., 2022), learning to prompt for continual learning (Wang et al., 2022), context optimization and conditional prompt learning for multi-modal models (Zhou et al., 2022b), and prompt-based domain adaptation strategies (Ge et al., 2023). For example, Nie et al. (2023) introduce the pro-tuning algorithm for learning task-specific vision prompts, applied to downstream task input images with the pre-trained model remaining frozen. Shen et al. (2024) leverage cross-task knowledge to optimize prompts, thereby enhancing the performance of vision-language models and avoiding the need to independently learn prompt vectors for each task from scratch. Cho et al. (2023) introduce distribution-aware prompt tuning for vision-language models, optimizing prompts by balancing inter-class dispersion and intra-class similarity. MaPLe (Khattak et al., 2023) further transfers text features to the visual encoder during prompt tuning to avoid overfitting. These approaches leverage learnable prompts to enhance model performance across various applications. Despite significant advancements in previous research, challenges remain in extracting semantic features from style-diversified images and optimizing templates within continuous prompt tuning. In this study, we employ both NLP and visual prompt tuning to optimize STEM educational content retrieval, enhancing retrieval accuracy and efficiency by adjusting prompt tokens.

B Motivation and Scenarios

In practical teaching scenarios, teachers often encounter the need for precise image retrieval, such as searching for hand-drawn sketches, student-created artistic images, blurry blackboard

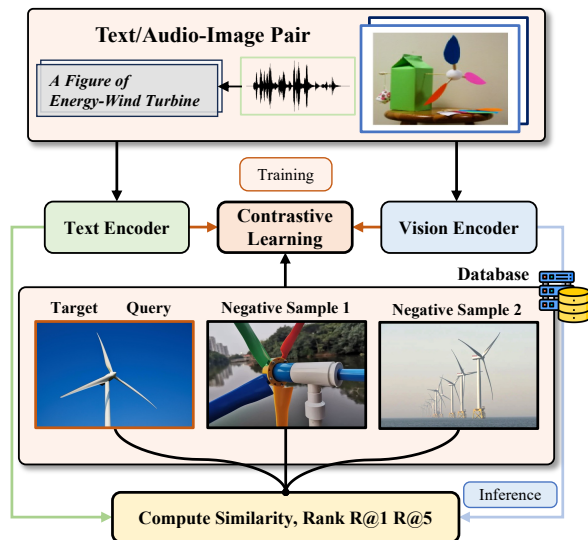


Figure 7: **The pipeline of Uni-Retrieval.** The image-text/audio pairs are input into their respective modality encoders. During the training procedure, contrastive learning is applied between the modality features of the positive samples (image-text/audio pairs) and the negative samples. During the inference procedure, the model calculates the similarity between the modality features of the query and the embeddings stored in the database. The retrieved results are ranked, and the performance is evaluated using $R@1/R@5$ as metrics.

drawings captured from a distance, classroom photographs of physical objects, or images from textbooks. However, current retrieval models predominantly focus on text-natural image queries, overlooking the diverse query styles common in educational contexts. This limitation makes it challenging for teachers to efficiently identify and retrieve educational images or texts tailored to diverse teaching scenarios, such as accurately setting learning contexts, articulating key teaching points, presenting instructional materials, and quickly locating supplementary resources.

Our proposed method enables teachers to query various styles’ answers with a range of retrieval approaches, including text, image, audio, or combinations of these modalities. This approach ensures fast and convenient retrieval, significantly reducing preparation time for teaching. Once teachers input their queries, our Uni-Retrieval system employs contrastive learning to compare images and text, calculating similarities based on attributes like objects, shapes, quantities, and orientations. The system ranks all database entries by similarity, outputting the top-1 or top-5 results to identify the most relevant and accurate teaching resource images, as illustrated

in Fig. 7. This approach empowers teachers to manage complex and dynamic teaching scenarios effortlessly, enhancing the clarity and effectiveness of STEM education.

C Experiments

In the database, the texts, their corresponding four images and audio structures for each dataset can share a single index, significantly reducing query time. All images and text are preprocessed using pretrained models to extract features, which are stored as embeddings. This approach eliminates the need for repeated feature extraction during use, saving time and reducing computational overhead, improving the efficiency of the retrieval system.

For the dataset selection, we choose four other datasets except our SER dataset, including the DSR dataset (Li et al., 2025), the ImageNet-X dataset, the SketchCOCO dataset (Gao et al., 2020a) and the DomainNet dataset (Peng et al., 2019). We use internVL-1.5 (Chen et al., 2024) to annotate the paint/sketch caption for the SketchCOCO and the DomainNet dataset. For the model in the prototype learning module, we choose the VGG (Simonyan, 2014) as the feature extractor. For the baseline selection, we apply two cross-modality pre-trained models (CLIP (Radford et al., 2021), BLIP (Li et al., 2023b)), two multi-modality pre-trained models (LanguageBind (Zhu et al., 2023), Unified-IO2 (Lu et al., 2024)), two style retrieval models (SceneTrilogy (Chowdhury et al., 2023b), FashionNTM (Pal et al., 2023)), four most recent cross-modality prompt learning models (VPT (Jia et al., 2022), CoCoOP (Zhou et al., 2022b), MaPLe (Khattak et al., 2023), FreestyleRet (Li et al., 2025)), and two database-driven retrieval models (GASKN (Yu et al., 2024), MKG (Zheng et al., 2024)) for the fair comparison. Specifically, we fine-tune the cross-modality models (CLIP, BLIP) on SER for convergence. We also train the prompt learning models on SER dataset based on VPT’s settings to adapt STEM style-diversified inputs. As for the multi-modality models, we evaluate the zero-shot performance on the style-diversified STEM education retrieval task due to multi-modality models’ comprehension ability on multi-style image inputs.

For the experiments on the SER dataset, Uni-Retrieval is initialized with OpenCLIP’s weights and trained on 8 A100 GPUs with batch size 24 per GPU and 20 training epochs. We use

AdamW as the optimizer, set the learning rate to $1e-5$ with a linearly warmed up operation in the first epochs and then decayed by the cosine learning rate schedule. The seed is set as 42. What’s more, all input images are resized into 224×224 resolution and then augmented by normalized operation. All text are padding zero to the max length of 20.

For the fine-tuning CLIP and BLIP models, all experiment settings are the same as Uni-Retrieval except the learning rate is set as $1e-6$. For prompt tuning models, we both use 4 prompt tokens to expand the token sequence. For all transformer-based models, we use the ViT-Large and 24-layers text transformer as the foundation models to keep balance between performance and efficiency.