CART: A Generative Cross-Modal Retrieval Framework with Coarse-To-Fine Semantic Modeling

Minghui Fang^{1*}, Shengpeng Ji^{1*}, Jialong Zuo¹, Hai Huang¹, Yan Xia¹, Jieming Zhu^{2†} Xize Cheng¹, Xiaoda Yang¹, Wenrui Liu¹, Gang Wang², Zhenhua Dong², Zhou Zhao^{1†}

¹Zhejiang University, ²Huawei Noah's Ark Lab minghuifang.zju@zju.edu.cn jiemingzhu@ieee.org zhaozhou@zju.edu.cn

Abstract

Cross-modal retrieval aims to search for instances, which are semantically related to the query through the interaction of different modal data. Traditional solutions utilize a singletower or dual-tower framework to explicitly compute the score between queries and candidates, which is challenged by training cost and inference latency with large-scale data. Inspired by the remarkable performance and efficiency of generative models, we propose a generative cross-modal retrieval framework (CART) based on coarse-to-fine semantic modeling, which assigns identifiers to each candidate and treats the generating identifier as the retrieval target. Specifically, we explore an effective coarse-to-fine scheme, combining K-Means and RQ-VAE to discretize multimodal data into token sequences that support autoregressive generation. Further, considering the lack of explicit interaction between queries and candidates, we propose a feature fusion strategy to align their semantics. Extensive experiments demonstrate the effectiveness of the strategies in the CART, achieving excellent results in both retrieval performance and efficiency.

1 Introduction

In recent years, multimedia data has exploded both in quantity and form, cross-modal retrieval (Zhu and Li, 2023; Huang et al., 2024; Liu et al., 2021; Xia et al., 2024) has become a research hot spot. Cross-modal retrieval tasks, which aim to retrieve relevant data from one modality (e.g. image (Radford et al., 2021; Cherti et al., 2023), audio (Elizalde et al., 2023a,b) or video (Luo et al., 2022; Wu et al., 2023)) based on a query from another modality (e.g. text (Wang et al., 2017)), play a crucial role in various multimedia applications.

Traditional cross-modal retrieval approaches are mainly based on single-tower or dual-tower frame-



Figure 1: Traditional single-tower / dual-tower retrieval matches the closest candidate to the query by calculating scores, while generative retrieval takes the generating candidate's identifier as the retrieval target.

work, as shown in Figure 1 (a). The single-tower model (Li et al., 2021, 2022, 2023a; Chen et al., 2024) performs fine-grained interactions between queries and candidates in a unified module, which shows excellent capability in retrieval performance, but is difficult to apply to large-scale corpora because of high latency. The dual-tower model (Radford et al., 2021; Luo et al., 2021; Elizalde et al., 2023b; Deshmukh et al., 2022; Girdhar et al., 2023; Zhu et al., 2023a) maps different modal data to the joint embedding space via two encoders, which improves the efficiency compared to the former. However, due to the modal gap, the dual-tower model often struggles to align multimodal data (Xia et al., 2024) effectively, thus slightly reducing accuracy.

Generative retrieval (Tay et al., 2022; Wang et al., 2022; Mehta et al., 2022; Zhuang et al., 2022; Ren et al., 2023; Li et al., 2023c) is an emerging paradigm in document retrieval that assigns identifiers (De Cao et al., 2020; Bevilacqua et al., 2022; Li et al., 2023c) to each candidate, followed by treating the generated identifiers as retrieval targets. This paradigm implements the retrieval process using a sequence-to-sequence model with excellent retrieval accuracy and efficiency, which sheds

^{*}Equal contribution.

[†]Corresponding authors.

new light on cross-modal retrieval. As shown in Figure 1 (b), there is no need to maintain a unified embedding space. Leveraging the power of generative models significantly enhances retrieval performance, and the generation speed remains independent of the dataset size.

While such generative framework has shown impressive performance in document retrieval, combining generative retrieval with cross-modal retrieval faces several challenges. Firstly, unlike documents where titles (Zhou et al., 2022b; Li et al., 2023c), strings (Bevilacqua et al., 2022; Tang et al., 2023), or keywords (Zhou et al., 2022a; De Cao et al., 2020; Zhang et al., 2023) can be directly used as identifiers, multimodal data are represented primarily through visual or auditory modalities. This means that we need effective discretization schemes to construct semantic identifiers for multimodal data, which is the core of generative retrieval. Secondly, semantic identifiers constructed from low-level visual or auditory information often exhibit a gap with the high-level semantics of natural language queries. Thirdly, generative retrieval encodes all candidates directly into the model parameters, lacking an explicit interaction process between queries and candidates.

In this work, we propose a generative retrieval framework, denoted as Cross-modal Autoregressive Retrieval Transformer (namely CART), designed to fully support end-to-end text-toimage/audio/video retrieval. The entire framework is divided into three modules: semantic identifier generation, caption enhancement, and feature fusion. Specifically, we propose combining the K-Means and RQ algorithms to generate identifiers with hierarchical semantic information for the candidates. Secondly, we perform caption generation for images/videos/audio, which serves as semantically aligned natural language queries. Finally, we design a two-branch coarse-to-fine feature fusion strategy, which is consistent with the hierarchical structure of semantic identifiers.

We conduct extensive experiments to demonstrate that CART has strong performance in text-toimage/audio/video retrieval. Meanwhile CART has stable retrieval speed on both CPU and GPU, meaning it is not affected by the number of candidates. Ablation studies and in-depth analyses validate the effectiveness and robustness of CART. CART enables end-to-end generative cross-modal retrieval based on a unified differentiable framework, providing a new solution for information retrieval in multimedia applications.

Our contributions are highlighted as follows.

- To the best of our knowledge, CART is the first generative cross-modal retrieval framework that comprehensively supports text-toimage/audio/video retrieval.
- We propose generating coarse-to-fine semantic identifiers, enabling a unified representation of multimodal data.
- We propose a coarse-fine feature fusion mechanism which effectively helps the retrieval model to perceive the query semantics.
- Extensive experiments demonstrate the superior performance of CART. Meanwhile, CART has stable retrieval efficiency and beats the single-tower model or dual-tower model under large-scale corpus.

2 Related Works

2.1 Cross-Modal Retrieval

Cross-modal retrieval aims to enable searches across different modalities. Currently, the mainstream approaches (Radford et al., 2021; Cherti et al., 2023; Elizalde et al., 2023a,b; Fang et al., 2021; Wu et al., 2023) to cross-modal retrieval tend to use either single-tower or dual-tower models, both of them formulate cross-modal retrieval as a discriminative problem, relying on discriminative loss and negative samples.

The difference is that the former performs finegrained interactions between the query and the candidates within a unified module, e.g., Cross-Attention, after which the score of each candidate is obtained by mapping in the output layer. BLIP (Li et al., 2022), BLIP-2 (Li et al., 2023a), InternVL-G (Chen et al., 2024) follow this setting to achieve high-accuracy retrieval in the domain of text-toimage retrieval; whereas, the latter tends to utilize two independent encoders to map different modalities into a joint feature space and then calculate their similarity by a distance function. CLIP (Radford et al., 2021), CLAP (Elizalde et al., 2023a) and CLIP2Video (Fang et al., 2021) follow this setting and complete the pre-training on large-scale datasets, achieving remarkable outcomes in textimage retrieval (Cao et al., 2022), text-audio retrieval (Koepke et al., 2022) and text-video retrieval (Zhu et al., 2023b), respectively.



Figure 2: A coarse-Fine semantic identifier generation strategy.

Both solutions have their own trade-offs; the single-tower model can achieve extremely high accuracy, but sacrifices efficiency because of finegrained interactions, leading it to be suitable only for small-scale retrieval tasks, and the dual-tower model effectively improves speed, but demands extensive data to train a unified representation space, leading to a slight decrease in retrieval accuracy. In contrast, CART aims to utilize generative models to improve retrieval efficiency while ensuring excellent retrieval performance.

2.2 Generative Retrieval

Generative retrieval (Tay et al., 2022; Wang et al., 2022; Long et al., 2024) represents an innovative paradigm in document retrieval. Its core idea is to construct indices for documents and replace the process of calculating relevance scores with generating identifiers.

The performance of generative retrieval is heavily influenced by candidate identifiers. Past studies have delved into the application of various identifier types across diverse scenarios, including keywords-based (Zhou et al., 2022b,a; De Cao et al., 2020; Zhang et al., 2023; Bian et al., 2025), Web URLs (Ren et al., 2023; Ziems et al., 2023), and substrings of paragraphs (Zhou et al., 2022b; Bevilacqua et al., 2022; Tang et al., 2023; Li et al., 2023c). Considering the absence of explicit document content, DSI (Tay et al., 2022) and NCI (Wang et al., 2022) recognize the challenge in generating document identifiers solely based on input queries. Therefore, they advocate using the hierarchical K-Means algorithm to inject prior knowledge into identifiers. In other words, documents sharing close semantic ties are assigned similar docids. This methodology seamlessly embeds the semantic information of documents into the decoding process, which facilitates the learning process

of the retrieval model. CART considers to build a unified identifier generation pipeline for multimodal data, combining K-Means and residual vector quantization to compress the embedding within a fixed step size. In addition, we investigate the application of generative retrieval techniques to recommendation tasks (Wang et al., 2024; Hong et al., 2025).

3 Method

3.1 Overview

Overall, CART is a framework that comprehensively generative text-tosupports image/video/audio retrieval, based on а coarse-to-fine semantic modeling scheme. The whole pipeline can be divided into three modules. (1) Semantic Identifier Construction: constructing semantically rich identifiers for all candidates based on the discretization method; (2) Caption Enhancement: generating captions for multimodal data as queries; (3) Feature **Fusion**: generative models are trained using the <query, identifier> pairs, and with the help of effective feature fusion mechanism to improve the performance. To enhance the robustness of CART, we incorporate consistency loss to suppress overfitting, and guarantee to generate valid identifiers via constrained beam search.

3.2 Coarse-Fine Identifier Generation

Images, audio, and video lack explicit attributes that can be directly used as identifiers, necessitating their discretization into token sequences in the latent space. The pipeline is illustrated in Figure 2.

3.2.1 Coarse Token

Intuitively, the first token of identifier is critical, if the first token of identifier is generated incorrectly, subsequent generation will be meaningless. With this in mind, we hope that the first token of the identifier captures the full semantic information of the item, so that the retrieval model can easily predict the semantically closest information based on a natural language query. We utilize K-Means (Hartigan and Wong, 1979) to cluster the embeddings of all the items, which are encoded from ImageBind (Girdhar et al., 2023), and the clustering result is considered as the first token of the semantic identifier. Although the K-Means algorithm can finish item categorization rapidly in the whole data space, it is difficult to consider the subtle gaps between items effectively. Therefore, we denote the clustering result k of K-Means as the coarse token.

3.2.2 Fine Token

Subsequently, we utilize original embeddings to subtract the embedding from the clustering center of the K-Means algorithm, in order to highlight the subtle differences between items. Then we use RQ-VAE (Zeghidour et al., 2021) to construct the remaining tokens in the semantic identifier.

RQ-VAE combines the advantages of variational autoencoder and residual quantization, where the autoencoder is jointly trained by updating the DNN encoder-decoder parameters and the quantization codebook. The encoder $E(\cdot)$ first downsamples the input x to learn the latent representation z = E(x), which removes noise and unimportant information and preserves the most meaningful features of the data. The quantizer $Q = \{C_1, C_2, \dots, C_M\}$ where M is the number of codebooks. The codebook $C_m = \{e_m^1, e_m^2, \dots, e_m^N\}$ where N is the codebook size and e_m^n denotes the *n*-th entry in the *m*-th codebook. The initial residual of the quantizer is defined as $r_0 = z$, then r_0 is quantized by mapping it to the nearest embedding e_1^k from the first codebook C_1 . The index of the closest embedding, i.e., $v_1 = argmin_k ||r_0 - e_1^k||$, represents the token of the semantic identifier obtained in codebook C_1 , the residual is updated to $r_1 = r_0 - e_1^{v_1}$. This process is executed recursively m times and we obtain the index list $V = \{v_1, v_2, \dots, v_M\}$ denoting the semantic identifier produced by the RQ-VAE. Note that all codebooks do not share parameters, they are updated independently to distinguish between the different semantics of each hierarchy. Finally $Q(z) = \sum_{i=1}^{M} e_i^{v_i}$ is passed to the decoder $D(\cdot)$ to try to reconstruct the input x. We train the RQ-VAE using two loss functions $\mathcal{L} = \mathcal{L}_{recon} + \alpha \mathcal{L}_{commit}$,

where:

$$\mathcal{L}_{recon} = \|x - D(Q(E(x)))\|_2^2,$$

$$\mathcal{L}_{commit} = \sum_{i=1}^M \|sg[r_i] - e_{v_i}^i\|_2^2 + \beta \|r_i - sg[e_{v_i}^i]\|_2^2$$

(1)

Here, both α and β are hyperparameters, and sg is the stop-gradient operation (Van Den Oord et al., 2017). RQ-VAE hierarchically computes the residuals of the features for each item. In the refinement process, it focuses on the variation of the features at different levels and can capture the subtle features of the items. Therefore, we denote the quantization result of RQ-VAE as the fine token.

3.2.3 Unique Token

In large-scale corpora, it is common to have semantically close items, which makes it difficult to avoid identifier collisions (Rajput et al., 2024). In order to assign unique identifiers to each item, we maintain a prefix database to detect conflicts. We perform a "search and count" operation on all identifiers, if the identifier already exists, increase its counter by one, on the contrary, insert the identifier into the prefix database with an initial counter value of 0. The value of the counter u will be appended to the end of the identifier is unique. Note that search, insert or update of the prefix database is so efficient that its time spent is even negligible.

Finally, we concatenate coarse token, fine token and unique token as identifiers for items, i.e. $(k, v_1, v_2, \dots, v_M, u)$.

3.3 Caption Enhancement

One huge challenge of generative retrieval is how to make the retrieval model aware of the semantic information represented by the identifiers. Since the content of each item is not explicitly known at inference, it must be incorporated into the model parameters during training, as indicated in previous work (Wang et al., 2022; Zhuang et al., 2022; Zhang et al., 2024).

To bridge this gap, we employ pre-trained multimodal models to generate captions for each candidate, taking information-rich captions as queries. The retrieval model is subsequently trained using <query, identifier> pairs, which effectively understands the information embedded in each token in the semantic identifier.



Figure 3: The architecture of the CART.

3.4 Coarse-to-Fine Feature Fusion

CART utilizes a standard encoder-decoder architecture, both made of stacks of attention layers. The encoder is in charge of capturing the semantics of the input query, the decoder reads the output of the encoder and generates semantic identifier token by token. Given an input query q, the output of an encoder with S layers can be defined as $E(q) = (E_1, E_2, \dots, E_S)$. Each encoder layer can capture different semantic representations (Cornia et al., 2020) of queries. The lower layer can capture some basic regional relevance, while the higher layer further refines and optimizes semantic representations based on the former. In light of this, we propose the coarse-to-fine feature fusion strategy. As shown in Figure 3, we process the encoder output in two branches.

3.4.1 Coarse Fusion

The outputs from each encoder layer are concatenated and passed through a fusion layer, to integrate information from different levels of the encoding process. Because of the "concatenate-then-fusion" method, we naturally call it coarse fusion, which is similar to the idea of coarse token. Formally,

$$Z = W[E_1, E_2, \dots, E_S] + b,$$
 (2)

where $[\cdot, \cdot]$ denotes connecting the outputs of all encoder layers from the last dimension, W and b are the learnable weights and bias, respectively. Subsequently, Z interacts with the decoder input Y through cross-attention and completes the postprocessing using self-gating. The whole process of coarse fusion $G(\cdot, \cdot)$ can be written as

$$C(Y, Z) = \text{Attention}(W_q Y, W_k Z, W_v Z),$$

$$G(Y, Z) = \frac{1}{S} \sigma(C(Y, Z)),$$
(3)

where $\mathcal{C}(\cdot, \cdot)$ stands for the encoder-decoder crossattention, S is the number of encoder layers and $\sigma(\cdot)$ is the sigmoid activation. W_q , W_k and W_v are matrices of learnable weights.

3.4.2 Fine Fusion

The purpose of fine fusion is for more careful deciding which subtle features to use. Intuitively, similar to mixture-of-experts (Shazeer et al., 2017), this design treats each encoder layer as an expert. The outputs of each layer in the encoder interact independently with the inputs of the decoder through cross-attention, while utilizing the weights α_i regulate each encoder layer's contribution and their relative importance. The decision-making process $L(\cdot, \cdot)$ can be written as

$$L(Y, E(q)) = \sum_{i=1}^{S} \alpha_i \odot \mathcal{C}(Y, E_i), \qquad (4)$$

and inspired by (Cornia et al., 2020), we introduce the memory-augmented mechanism for weight computation, as follows

$$\alpha_i = \sigma(W_i[Y, \mathcal{C}(Y, E_i)] + b), \tag{5}$$

where $[\cdot, \cdot]$ indicates concatenation, $\sigma(\cdot)$ is the sigmoid activation, W_i is the learnable weight matrices and b is bias.

Finally, we add the coarse feature and fine feature as input to the next decoder layer.

3.5 Training and Inference

We use <query, identifier> pairs to train a sequenceto-sequence model. Given an input query q, the probability of generating the semantic identifier Tcan be written:

$$p(T \mid E(q), \theta) = \prod_{i=1}^{J} p(t_i \mid E(q), t_{< i}, \theta_i), \quad (6)$$

where t_i is the *i*-th token in the semantic identifier T, J denotes the length of T and $E(\cdot)$ denotes the encoder, θ denotes the total parameters and θ_i is the parameter for the *i*-th step.

To mitigate the overfitting problem, we introduce a bidirectional KL divergence loss (Wu et al., 2021) to train the decoding process. For the user query, we denote the decoder representations by two forward passes with independent dropouts before softmax as P and Q, i.e.,

$$\mathcal{L}_{KL} = \sum_{i} P(i) \log\left(\frac{P(i)}{Q(i)}\right) + \sum_{i} Q(i) \log\left(\frac{Q(i)}{P(i)}\right)$$
(7)

As shown in Figure 3, given a collection of training examples $\mathcal{D} = \{(q, d)\}$ composed of queries (training queries and augmented queries) and identifiers, the loss function can be written as follows:

$$\mathcal{L}(\theta) = \sum_{(q,d)\in\mathcal{D}} (logp(d \mid E(q), \theta) + \omega \mathcal{L}_{KL}),$$
(8)

where $p(d \mid E(q), \theta)$ denotes the probability of generating d with q as the input and ω denotes a scaling factor of KL divergence loss.

In the inference stage, we execute a constrained beam search on the decoder network. Benefiting from the prefix database constructed when generating identifiers, we can effectively build a prefix tree for all identifiers, which will restrict the model to generating only valid identifiers. Due to the hierarchical nature of identifiers, it is convincing (Wang et al., 2022) to constrain the beam search decoding process with a prefix tree.

4 **Experiments**

We evaluate the performance of CART on text-toimage/audio/video, respectively.

4.1 Datasets & Baselines

Datasets We selected six widely used datasets: Flickr30K (Young et al., 2014), MS-COCO (Lin et al., 2014), Clotho (Drossos et al., 2020), AudioCaps (Kim et al., 2019), MSR-VTT (Xu et al.,

Table 1: Compared to single-tower models, results from original papers.

Mathod		Flickr30	K	1	MS-COC	0	Throughput
Method	R@1	R@5	R@10	R@1	R@5	R@10	Throughput
BLIP-2	89.7	98.1	98.9	66.3	86.5	91.8	1.68/s
InternVL-G	85.0	97.0	98.6	58.6	81.3	88.0	2.03/s
CART(Ours)	81.8	96.1	98.4	52.4	77.5	86.1	105.8/s

2016) and MSVD (Chen and Dolan, 2011). Further details are provided in Appendix A.

Baselines We selected single-tower, dual-tower, and generative retrieval architectures as baselines. Further details are provided in Appendix B.

4.2 Implementation Details

The implementation mainly consists of the semantic identifier generation pipeline and the retrieval pipeline, with details on network structure and parameter settings provided in Appendix C. All experiments are based on 4 NVIDIA V100 GPUs with training batch size set to 256.

4.3 Empirical Results

4.3.1 Metrics

Consistent with prior studies, we use widely accepted metrics for information retrieval, including Recall@K and Mean Reciprocal Rank (MRR). The details about the metrics are in Appendix D.

4.3.2 Comparison of single-tower retrieval

As shown in Table 1, We observe that the singletower model has an overwhelming advantage on R@1, thanks to the fine-grained interactions between queries and candidates. Notably, with the increase in the number of recalls, the gap between CART and BLIP-2 (Li et al., 2023a) narrows faster than that between InternVL-G (Chen et al., 2024) and BLIP-2, which proves that CART has comparable performance. Afterwards, we simulated 1M candidate images and experimented with 100 concurrent queries. Table 1 shows the huge latency of the single-tower model, which is not applicable to large-scale retrieval. In contrast, CART achieves a better balance between performance and efficiency.

4.3.3 Comparison of dual-tower retrieval

As shown in in Table 2, we report the retrieval results for CART and the corresponding baselines. In text-to-image/video/audio retrieval, CART exhibits strong performance, leading the baseline in all metrics. Compared to images, videos and audios are temporal and thus contain more semantic

Taalr	Mathad		F	lickr30k			М	S-COCO		
Täsk	Method	R@1	R@5	R@10	MRR@10	R@1	R@5	R@10	MRR@10	
	CLIP	55.9	82.8	90.6	67.3	30.4	55.9	66.8	41.3	
	OpenCLIP	63.9	87.3	93.2	74.1	39.4	65.4	75.6	50.4	
	MobileCLIP	74.9	92.9	96.3	82.6	50.3	74.9	82.4	60.8	
Text-to-Image Retrieval	ONE-PEACE(Pretrained)	73.4	91.5	95.4	81.2	48.0	71.6	79.5	58.0	
-	ImageBind(Huge)	74.9	93.0	96.1	82.7	49.4	73.3	81.5	59.6	
	LanguageBind(Image)	69.5	90.8	94.9	78.7	45.3	69.8	78.6	55.8	
	CART(ours)	81.8	96.1	98.4	88.0	52.4	77.5	86.1	63.2	
T1-	M-41 J	Clotho			AudioCaps					
Task	Method	R@1	R@5	R@10	MRR@10	R@1	R@5	R@10	MRR@10	
	CLAP-HTSAT	16.1	38.3	51.1	25.7	36.1	71.8	83.9	52.1	
	HTSAT-BERT	19.7	45.7	59.4	37.6	42.2	76.5	87.1	57.4	
Text-to-Audio Retrieval	LanguageBind(Audio)	16.3	40.4	53.9	32.9	15.2	49.8	67.6	29.8	
	VALOR-B	17.5	42.7	55.3	34.5	40.1	73.9	83.1	58.7	
	ONE-PEACE(Pretrained)	22.4	49.0	62.7	40.7	42.5	77.5	88.4	60.5	
	CART(ours)	46.4	70.6	76.0	52.3	49.8	79.8	86.8	62.0	
	M - 41 J	MSR-VTT					MSVD			
Task	Method	R@1	R@5	R@10	MRR@10	R@1	R@5	R@10	MRR@10	
	CLIP2Video	45.6	72.5	81.7	-	47.0	76.8	85.9	-	
	CLIP4Clip	44.5	71.4	81.6	-	46.2	76.1	84.6	-	
	UMT-B	46.3	72.7	82.0	-	47.4	76.8	84.0	-	
Text-to-Video Retrieval	Cap4Video	49.3	74.3	83.8	-	51.8	80.8	88.3	-	
	ImageBind(Huge)	36.8	61.8	70.0	46.4	39.3	67.5	78.5	53.4	
	LanguageBind(Video)	42.7	67.1	77.0	54.4	53.5	80.5	87.5	60.6	
	CART(ours)	52.6	75.4	84.2	61.8	63.6	87.9	92.8	73.6	

Table 2: Compared to dual-tower models, the results are from the original papers or official reproduction.

Table 3: Compared to generative retrieval models.

Douodiam Mathad]	Flickr30	К	Ν	MS-COCO			
Faradigin	Wethod	R@1	R@5	R@10	R@1	R@5	R@10		
	Semantic ID	22.9	34.9	37.4	13.3	30.4	35.9		
GRACE	Structured ID	37.4	59.5	66.2	16.7	39.2	50.3		
	Atomic ID	68.4	88.9	93.7	41.5	69.1	79.1		
CART(Ours)	Coarse-Fine ID	81.78	96.14	98.38	52.42	77.53	86.11		

information, which increases the difficulty of retrieval. While CART achieves promising results benefiting from coarse-fine semantic identifiers and effective training strategies, proving the feasibility of generative retrieval. We also performed a more detailed efficiency comparison with the dual-tower model, as shown in Figure 4.

4.3.4 Comparison of generative retrieval

As shown in Table 3, GRACE (Li et al., 2024) performs relatively poorly because the multiple identifier schemes it employs are predefined and do not comprehensively consider semantic information, whereas CART employs a novel coarse-to-fine-grained identifier generation pipeline to effectively facilitate model learning. It is worth noting that GRACE adopts Flamingo (Alayrac et al., 2022) (a MLLM) as the model backbone, while CART adopts a naïve small-parameter transformer, the final difference in retrieval results is enough to prove the importance of identifier for generative retrieval.

4.3.5 Ablation experiment

Furthermore, to investigate the effect of each component, we report the ablation results in Table 4. The detailed analysis is described below.

w/o consistency loss This setting removes the consistency-based regularization loss, and we observe that the model is more prone to overfitting, which leads to degradation of performance.

w/o fusion strategy This setting removes the process of feature fusion in the decoder, demonstrating the importance of fully considering outputs from all encoder layers.

w/o K-Means This setting removes the process of K-Means and uses only RQ-VAE to construct semantic identifiers. Experiments show that clustering the original embeddings via K-Means brings prior knowledge that is effective for retrieval.

w/o RQ-VAE This setting removes the process of RQ-VAE and uses only the hierarchical K-Means to construct K-fork trees for the original embeddings, where the paths from the root node to the leaf nodes are considered as identifiers of the items. Experiments show that hierarchical K-Means loses semantic information between different clusters, which is consistent with Rajput et al. (2024).

Table 4: Ablation Study on Flickr30k, we leave the other ablation results in Appendix E.

Setting	R@1	R@5	R@10	MRR@10
w/o consistency loss	81.64	95.94	98.04	87.85
w/o fusion strategy	75.54	93.54	96.72	83.11
w/o K-Means	79.50	94.92	97.52	86.12
w/o RQ-VAE	76.22	92.86	96.16	83.31
CART	81.78	96.14	98.38	88.04



Figure 4: The efficiency of CLIP, CLAP and CART are measured by throughput (queries processed per second).

4.4 In-depth Analysis

Efficiency Analysis We evaluate the retrieval efficiency of CLIP (Radford et al., 2021), CLAP (Elizalde et al., 2023b), and CART on both a CPU and an NVIDIA V100 GPU, simulating 100 concurrent queries while varying the number of candidates. The metric is throughput (Queries/per second) and the detailed results are shown in Figure 4. CLIP and CLAP pre-encode candidates into embeddings, where the primary computational costs arise from text encoding, similarity computation, and ranking. In contrast, CART employs beam search to generate identifiers, integrating retrieval and ranking. As the number of candidates increases, the efficiency of dual-tower models like CLIP and CLAP declines due to increased computational demands. However, CART maintains stable efficiency on both GPU and CPU, with its advantage becoming more pronounced as candidate numbers grow, attributed to encoding all items directly as parameters.

Feature Fusion Configuration The effectiveness of coarse-to-fine feature fusion is influenced by the number of encoder layers. We evaluate configurations with {1,2,3,4,5} layers, with results summarized in Table 5. The metrics show consistent performance improvement as the number of layers increases, validating the fusion approach. However, excessive layers may lead to overfit-



Figure 5: The t-SNE visualization of item embeddings which have the same token prefixes.

Table 5: Results with varying encoder layers.

Layers	R@1	R@5	R@10	MRR@10
1	76.16	94.48	97.78	84.11
2	80.52	95.16	97.92	86.87
3	81.78	96.14	98.38	88.04
4	82.10	95.68	98.22	88.11
5	81.08	95.64	97.78	87.37

ting and performance degradation. Consequently, CART employs a 3-layer encoder for optimal balance. Correspondingly, we provide an analysis of the decoder layer in Appendix F.

Distribution of Semantic Identifiers We analyzed the distribution of candidate embeddings with the same token prefix. Figure 5 shows tighter intra-cluster distributions and more dispersed intercluster distributions, indicating that the identifiers capture the semantics of the candidates.

Identifier Configuration We investigate the impact of semantic identifier length and size through experiments on Flickr30k. Results in Tables 6 and 7 highlight the robustness of CART, with further analysis detailed in Appendix G. We also provide an analysis of the codebook vector dimensions in Appendix H.

5 Conclusion

In this work, we propose CART, a novel generative cross-modal retrieval framework that can fully support text-to-image/video/audio retrieval. Benefiting from the design of coarse-fine semantic identifiers and feature fusion strategy, CART can effectively understand user queries and accurately retrieve candidates. We conducted extensive experiments to demonstrate the effectiveness and robustness of the framework, along with in-depth analysis for a comprehensive understanding of CART. Compared to single-tower and dual-tower frameworks, CART achieves a fine balance between retrieval performance and efficiency. We sincerely hope this will inspire future work.

Κ	Codebook	R@1	R@5	R@10	MRR@10
64	64*64	75.76	92.91	96.48	83.12
	128*128	81.22	95.52	97.54	87.39
	256*256	79.84	95.58	98.18	86.44
128	64*64	78.64	93.96	97.08	85.31
	128*128	81.78	96.14	98.38	88.04
	256*256	80.86	95.61	98.01	87.26
256	64*64	77.68	94.81	97.62	84.93
	128*128	81.54	95.81	98.31	87.68
	256*256	80.48	95.96	98.34	87.21

Table 6: Results with different identifier sizes.

Table 7: Results with different identifier lengths.

Length	R@1	R@5	R@10	MRR@10
4 5	81.78 81.84	96.14 96.38	98.38 98.44	88.04 88.51
6	80.54	95.44	98.06	87.57

Limitation

We initially explored the potential of generative models to accomplish cross-modal retrieval. However, in reality corpora are usually dynamically changing, and we would like to introduce continuous learning to accommodate this setting in our next work. Meanwhile, the current version only supports natural language queries, and we would like to explore a unified embedding space to accommodate multimodal data queries.

Acknowledgement

We thank MindSpore (https://www.mindspore. cn) for the partial support of this work, which is a new deep learning computing framework.

References

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736.
- Michele Bevilacqua, Giuseppe Ottaviano, Patrick Lewis, Scott Yih, Sebastian Riedel, and Fabio Petroni. 2022. Autoregressive search engines: Generating substrings as document identifiers. Advances in Neural Information Processing Systems, 35:31668–31683.
- Zhipeng Bian, Jieming Zhu, Xuyang Xie, Quanyu Dai, Zhou Zhao, and Zhenhua Dong. 2025. Mira: Empowering one-touch ai services on smartphones with mllm-based instruction recommendation. In *ACL*.

- Min Cao, Shiping Li, Juntao Li, Liqiang Nie, and Min Zhang. 2022. Image-text retrieval: A survey on recent research and development. *arXiv preprint arXiv:2203.14713*.
- David Chen and William B Dolan. 2011. Collecting highly parallel data for paraphrase evaluation. In *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, pages 190–200.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, 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, pages 24185–24198.
- Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. 2023. Reproducible scaling laws for contrastive language-image learning. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 2818–2829.
- Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. 2020. Meshed-memory transformer for image captioning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10578–10587.
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2020. Autoregressive entity retrieval. *arXiv preprint arXiv:2010.00904*.
- Soham Deshmukh, Benjamin Elizalde, and Huaming Wang. 2022. Audio retrieval with wavtext5k and clap training. *arXiv preprint arXiv:2209.14275*.
- Konstantinos Drossos, Samuel Lipping, and Tuomas Virtanen. 2020. Clotho: An audio captioning dataset. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing* (*ICASSP*), pages 736–740. IEEE.
- Benjamin Elizalde, Soham Deshmukh, Mahmoud Al Ismail, and Huaming Wang. 2023a. Clap learning audio concepts from natural language supervision. In ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 1–5. IEEE.
- Benjamin Elizalde, Soham Deshmukh, and Huaming Wang. 2023b. Natural language supervision for general-purpose audio representations. *Preprint*, arXiv:2309.05767.
- Han Fang, Pengfei Xiong, Luhui Xu, and Yu Chen. 2021. Clip2video: Mastering video-text retrieval via image clip. *arXiv preprint arXiv:2106.11097*.
- Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. 2020. Multi-modal transformer for video retrieval. In *Computer Vision–ECCV 2020:*

16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16, pages 214–229. Springer.

- Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. 2023. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15180–15190.
- John A Hartigan and Manchek A Wong. 1979. Algorithm as 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)*, 28(1):100–108.
- Minjie Hong, Yan Xia, Zehan Wang, Jieming Zhu, Ye Wang, Sihang Cai, Xiaoda Yang, Quanyu Dai, Zhenhua Dong, Zhimeng Zhang, and Zhou Zhao. 2025. EAGER-LLM: enhancing large language models as recommenders through exogenous behaviorsemantic integration. In WWW, pages 2754–2762.
- Hai Huang, Yan Xia, Shengpeng Ji, Shulei Wang, Hanting Wang, Jieming Zhu, Zhenhua Dong, and Zhou Zhao. 2024. Unlocking the potential of multimodal unified discrete representation through training-free codebook optimization and hierarchical alignment. *arXiv preprint arXiv:2403.05168*.
- Chris Dongjoo Kim, Byeongchang Kim, Hyunmin Lee, and Gunhee Kim. 2019. Audiocaps: Generating captions for audios in the wild. In *Proceedings of the* 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pages 119–132.
- A Sophia Koepke, Andreea-Maria Oncescu, Joao Henriques, Zeynep Akata, and Samuel Albanie. 2022. Audio retrieval with natural language queries: A benchmark study. *IEEE Transactions on Multimedia*.
- Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. 2018. Stacked cross attention for image-text matching. In *Proceedings of the European conference on computer vision (ECCV)*, pages 201–216.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. Blip-2: Bootstrapping language-image pretraining with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pretraining for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. Advances in neural information processing systems, 34:9694–9705.

- Kunchang Li, Yali Wang, Yizhuo Li, Yi Wang, Yinan He, Limin Wang, and Yu Qiao. 2023b. Unmasked teacher: Towards training-efficient video foundation models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19948–19960.
- Yongqi Li, Wenjie Wang, Leigang Qu, Liqiang Nie, Wenjie Li, and Tat-Seng Chua. 2024. Generative cross-modal retrieval: Memorizing images in multimodal language models for retrieval and beyond. *arXiv preprint arXiv:2402.10805*.
- Yongqi Li, Nan Yang, Liang Wang, Furu Wei, and Wenjie Li. 2023c. Multiview identifiers enhanced generative retrieval. arXiv preprint arXiv:2305.16675.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Computer Vision– ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer.
- Jing Liu, Sihan Chen, Xingjian He, Longteng Guo, Xinxin Zhu, Weining Wang, and Jinhui Tang. 2024. Valor: Vision-audio-language omni-perception pretraining model and dataset. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Jing Liu, Xinxin Zhu, Fei Liu, Longteng Guo, Zijia Zhao, Mingzhen Sun, Weining Wang, Hanqing Lu, Shiyu Zhou, Jiajun Zhang, et al. 2021. Opt: Omniperception pre-trainer for cross-modal understanding and generation. *arXiv preprint arXiv:2107.00249*.
- Xinwei Long, Jiali Zeng, Fandong Meng, Zhiyuan Ma, Kaiyan Zhang, Bowen Zhou, and Jie Zhou. 2024. Generative multi-modal knowledge retrieval with large language models. *arXiv preprint arXiv:2401.08206*.
- Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. 2021. CLIP4Clip: An empirical study of clip for end to end video clip retrieval. *arXiv preprint arXiv:2104.08860*.
- Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. 2022. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508:293– 304.
- Sanket Vaibhav Mehta, Jai Gupta, Yi Tay, Mostafa Dehghani, Vinh Q Tran, Jinfeng Rao, Marc Najork, Emma Strubell, and Donald Metzler. 2022. Dsi++: Updating transformer memory with new documents. *arXiv preprint arXiv:2212.09744*.
- Xinhao Mei, Chutong Meng, Haohe Liu, Qiuqiang Kong, Tom Ko, Chengqi Zhao, Mark D Plumbley, Yuexian Zou, and Wenwu Wang. 2024. Wavcaps: A chatgpt-assisted weakly-labelled audio captioning dataset for audio-language multimodal research. *IEEE/ACM Transactions on Audio, Speech, and Language Processing.*

- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. *arXiv preprint arXiv:1904.01038*.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Shashank Rajput, Nikhil Mehta, Anima Singh, Raghunandan Hulikal Keshavan, Trung Vu, Lukasz Heldt, Lichan Hong, Yi Tay, Vinh Tran, Jonah Samost, et al. 2024. Recommender systems with generative retrieval. *Advances in Neural Information Processing Systems*, 36.
- Ruiyang Ren, Wayne Xin Zhao, Jing Liu, Hua Wu, Ji-Rong Wen, and Haifeng Wang. 2023. Tome: A two-stage approach for model-based retrieval. *arXiv* preprint arXiv:2305.11161.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Yubao Tang, Ruqing Zhang, Jiafeng Guo, Jiangui Chen, Zuowei Zhu, Shuaiqiang Wang, Dawei Yin, and Xueqi Cheng. 2023. Semantic-enhanced differentiable search index inspired by learning strategies. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4904–4913.
- Yi Tay, Vinh Tran, Mostafa Dehghani, Jianmo Ni, Dara Bahri, Harsh Mehta, Zhen Qin, Kai Hui, Zhe Zhao, Jai Gupta, et al. 2022. Transformer memory as a differentiable search index. *Advances in Neural Information Processing Systems*, 35:21831–21843.
- Aaron Van Den Oord, Oriol Vinyals, et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Pavan Kumar Anasosalu Vasu, Hadi Pouransari, Fartash Faghri, Raviteja Vemulapalli, and Oncel Tuzel. 2023. Mobileclip: Fast image-text models through multi-modal reinforced training. *arXiv preprint arXiv:2311.17049*.
- Bokun Wang, Yang Yang, Xing Xu, Alan Hanjalic, and Heng Tao Shen. 2017. Adversarial cross-modal retrieval. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 154–162.

- Peng Wang, Shijie Wang, Junyang Lin, Shuai Bai, Xiaohuan Zhou, Jingren Zhou, Xinggang Wang, and Chang Zhou. 2023. One-peace: Exploring one general representation model toward unlimited modalities. *arXiv preprint arXiv:2305.11172*.
- Ye Wang, Jiahao Xun, Minjie Hong, Jieming Zhu, Tao Jin, Wang Lin, Haoyuan Li, Linjun Li, Yan Xia, Zhou Zhao, and Zhenhua Dong. 2024. EAGER: two-stream generative recommender with behaviorsemantic collaboration. In *KDD*, pages 3245–3254.
- Yujing Wang, Yingyan Hou, Haonan Wang, Ziming Miao, Shibin Wu, Qi Chen, Yuqing Xia, Chengmin Chi, Guoshuai Zhao, Zheng Liu, et al. 2022. A neural corpus indexer for document retrieval. *Advances in Neural Information Processing Systems*, 35:25600– 25614.
- Lijun Wu, Juntao Li, Yue Wang, Qi Meng, Tao Qin, Wei Chen, Min Zhang, Tie-Yan Liu, et al. 2021. R-drop: Regularized dropout for neural networks. *Advances in Neural Information Processing Systems*, 34:10890– 10905.
- Wenhao Wu, Haipeng Luo, Bo Fang, Jingdong Wang, and Wanli Ouyang. 2023. Cap4video: What can auxiliary captions do for text-video retrieval? In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 10704–10713.
- Yan Xia, Hai Huang, Jieming Zhu, and Zhou Zhao. 2024. Achieving cross modal generalization with multimodal unified representation. *Advances in Neural Information Processing Systems*, 36.
- Jun Xu, Tao Mei, Ting Yao, and Yong Rui. 2016. Msrvtt: A large video description dataset for bridging video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5288–5296.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78.
- Youngjae Yu, Jongseok Kim, and Gunhee Kim. 2018. A joint sequence fusion model for video question answering and retrieval. In *Proceedings of the European conference on computer vision (ECCV)*, pages 471–487.
- Neil Zeghidour, Alejandro Luebs, Ahmed Omran, Jan Skoglund, and Marco Tagliasacchi. 2021. Soundstream: An end-to-end neural audio codec. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30:495–507.
- Hailin Zhang, Yujing Wang, Qi Chen, Ruiheng Chang, Ting Zhang, Ziming Miao, Yingyan Hou, Yang Ding, Xupeng Miao, Haonan Wang, et al. 2024. Modelenhanced vector index. Advances in Neural Information Processing Systems, 36.

- Peitian Zhang, Zheng Liu, Yujia Zhou, Zhicheng Dou, and Zhao Cao. 2023. Term-sets can be strong document identifiers for auto-regressive search engines. *arXiv preprint arXiv:2305.13859*.
- Yujia Zhou, Jing Yao, Zhicheng Dou, Ledell Wu, and Ji-Rong Wen. 2022a. Dynamicretriever: A pre-training model-based ir system with neither sparse nor dense index. *arXiv preprint arXiv:2203.00537*.
- Yujia Zhou, Jing Yao, Zhicheng Dou, Ledell Wu, Peitian Zhang, and Ji-Rong Wen. 2022b. Ultron: An ultimate retriever on corpus with a model-based indexer. *arXiv preprint arXiv:2208.09257*.
- Bin Zhu, Bin Lin, Munan Ning, Yang Yan, Jiaxi Cui, HongFa Wang, Yatian Pang, Wenhao Jiang, Junwu Zhang, Zongwei Li, et al. 2023a. Languagebind: Extending video-language pretraining to n-modality by language-based semantic alignment. *arXiv preprint arXiv:2310.01852*.
- Cunjuan Zhu, Qi Jia, Wei Chen, Yanming Guo, and Yu Liu. 2023b. Deep learning for video-text retrieval: a review. *International Journal of Multimedia Information Retrieval*, 12(1):3.
- Yi Zhu and Xiu Li. 2023. Iterative uni-modal and crossmodal clustered contrastive learning for image-text retrieval. In 2023 International Conference on Pattern Recognition, Machine Vision and Intelligent Algorithms (PRMVIA), pages 15–23.
- Shengyao Zhuang, Houxing Ren, Linjun Shou, Jian Pei, Ming Gong, Guido Zuccon, and Daxin Jiang. 2022. Bridging the gap between indexing and retrieval for differentiable search index with query generation. *arXiv preprint arXiv:2206.10128*.
- Noah Ziems, Wenhao Yu, Zhihan Zhang, and Meng Jiang. 2023. Large language models are builtin autoregressive search engines. *arXiv preprint arXiv:2305.09612*.

A Datasets

Text to Image Retrieval Flickr30K (Young et al., 2014) contains 31,783 images and each image is associated with 5 human-annotated sentences. We select 29,783 images for training, 1000 images for validation and 1000 images for testing. MS-COCO (Lin et al., 2014) comprises 123,287 images, and each image comes with 5 sentences of annotations. We followed the data split proposed in (Lee et al., 2018) and utilized 113,287 images for training, 5000 images for validation and 5000 images for testing.

Text to Audio Retrieval Each audio in Clotho (Drossos et al., 2020) has 5 manually annotated captions, and there are 3839 audios for training, 1045 audios for validation, and 1045 audios for

testing. Since some of the videos in YouTube are no longer available, we do not have the full collection of AudioCaps (Kim et al., 2019). Finally, we have 37,869 audios for training, 384 audios for validation, and 737 audios for testing.

Text to Video Retrieval MSR-VTT (Xu et al., 2016) is a dataset composed of 10,000 videos, each with a length that ranges from 10 to 32 seconds and 200,000 captions. We use 9,000 videos for training that follows the data split of (Gabeur et al., 2020), and the test set consists of 1,000 video-text pairs from (Yu et al., 2018). MSVD (Chen and Dolan, 2011) contains 1,970 videos, each with a length that ranges from one to 62 seconds. Train, validation and test splits contain 1,200, 100, and 670 videos, respectively. This is an appendix.

B Baselines

The single-tower model has powerful performance but high latency, which is not suitable for largescale retrieval, so we only select BLIP-2 (Li et al., 2023a) and InternVL-G (Chen et al., 2024) as baseline. We focus on comparing the dual-tower models, including the classic CLIP (Radford et al., 2021), as well as the iteratively upgraded Open-CLIP (Cherti et al., 2023), MobileCLIP (Vasu et al., 2023), CLAP (Elizalde et al., 2023b; Mei et al., 2024), VALOR-B (Liu et al., 2024), CLIP2Video (Fang et al., 2021), CLIP4Clip (Luo et al., 2022), UMT-B (Li et al., 2023b), Cap4Video (Wu et al., 2023), ImageBind (Girdhar et al., 2023), Language-Bind (Zhu et al., 2023a) and ONE-PEACE (Wang et al., 2023), which are utilized to perform textto-picture/video/audio retrieval, respectively. Furthermore, we note that GRACE (Li et al., 2024) enables generative text-to-image retrieval, so we add it to the baseline for generative cross-modal retrieval.

C Implementation Details

Semantic Identifier Generation Pipeline We use ImageBind to unify the processing of multimodal data into a latent space, ultimately obtaining 1024-dimensional embeddings. We employ the default mini batch K-Means algorithm in scikit-learn (Pedregosa et al., 2011), where K = 128, to cluster 1024-dimensional item embeddings. The encoder in RQ-VAE has three intermediate layers of size 512, 256 and 128 with ELU activation, and final latent representation dimension of 64, with which the deocoder is completely symmetric. There are 2 levels of codebook in the quantizer, and for each level, a codebook with cardinality 128 is maintained, where each vector in the codebook has a dimension of 64. Because we use a prefix database to ensure that each identifier is unique, this means at least $128^3 = 2,097,152$ items can be represented. We train the RQ-VAE using the Adam optimizer with an initial learning rate of 1e-6 and Inverse Square Root scheduler integrated in fairseq (Ott et al., 2019). The learning rate increases to 1e-4 after 300 warm-up epochs, and then gradually decreases until 500 epochs.

Retrieval Pipeline We initialize all parameters with Xavier uniform distribution. We use Adam optimizer with an initial learning rate of 1e-6, while increasing the learning rate to 1e-4 in 5 epochs using the Cosine Annealing Warmup Restarts scheduler. We set the scaling factor of the consistency-based regularization loss as $\omega = 0.15$.

D Metrics

Recall@K measures how often the desired item is hit in the top K retrieved candidates, where we set K to 1, 5 and 10 in experiments. MRR calculates the reciprocal of the rank at which the first relevant item is retrieved. A high MRR indicates that the relevant item has a high ranking position.

E Ablation Study

As shown in Table 8 and 9, we demonstrate the effectiveness of each module in CART on the text-to-image/audio/video retrieval task.

When **w/o consistency**, the model undergoes overfitting which also leads to degradation of re-trieval performance.

When **w/o fusion strategy**, the metrics show different degrees of decrease, proving the effectiveness of the coarse-to-fine feature fusion mechanism.

When w/o K-Means token or w/o RQ-VAE token, the retrieval performance also decreases, which proves the effectiveness of coarse-fine token.

F Decoder Layers

In the previous section, we perform ablation experiments on the number of encoder layers, and correspondingly, we perform ablation experiments on the number of decoder layers.

As shown in Table 10, as the number of decoder layers increases, the capability of the model gradu-

ally increases and the retrieval performance gradually improves, but too many decoder layers can lead to the model undergoing serious overfitting, leading to performance degradation.

G Identifier Configuration

As shown in Table 6 We try to vary the K in K-Means as well as the codebook size in RQ-VAE, and observe that the retrieval performance of the model is more robust for most settings. Only when K is set to 64 or codebook size is set to 64, the model performance shows a more significant degradation. We analyze that too small a cluster category will lead to lower differentiation of K-Means and fail to achieve good clustering effect, and too small codebook size will likewise lead to limited semantic differentiation ability, resulting in larger information loss. Furthermore, it will also lead to an increase in the amount of items with the same identifiers, thus the model can only rely on the unique token which do not have any semantic information. We try to add codebook in RQ-VAE to increase the length of semantic identifiers and we observe that the retrieval performance fluctuates only slightly, as shown in Table 7. This demonstrates the effectiveness of the semantic identifiers and the stability of the retrieval model. Considering that too long semantic identifiers lead to lower codebook utilization and increased decoding steps, we only use four tokens. It is worth noting that retrieval performance exceeds the baseline in all settings, despite varying the length and size of the identifiers.

H Codebook Dimension

We further explored the codebook dimension by keeping the number of clustering centers at 128 and conducting experiments on Flickr30k, as shown in Table 11. We note that large dimensions not only make the samples sparse and affect the distance discrimination, but also may introduce more noise and affect the model learning semantics. And small dimensions will reduce the learning ability of codebook and make it difficult to revert the data features. Note that these parameter settings are beneficial for this experiment, and there may be some mathematical connection between them, which future work could explore more to help determine the best parameters for different datasets.

Satting		(Clotho				Au	dioCaps	
Setting	R@1	R@5	R@10	MRR@10		R@1	R@5	R@10	MRR@10
w/o consistency loss	44.51	68.76	74.56	49.87	2	47.86	76.74	83.89	57.68
w/o fusion strategy	43.44	66.15	73.03	48.92	2	46.04	76.96	82.21	55.01
w/o K-Means	45.81	69.31	74.05	51.59	2	47.21	77.11	85.19	52.88
w/o RQ-VAE	44.39	68.91	75.11	49.51	4	44.94	73.27	82.45	50.34
CART	46.4	70.6	76.0	52.3		49.8	79.8	86.8	62.0
Satting		M	SR-VTT				Ν	ASVD	
Setting	R@1	R@5	R@10	MRR@10]	R@1	R@5	R@10	MRR@10
w/o consistency loss	51.5	75.1	82.7	61.53	(61.19	87.16	92.38	72.03
w/o fusion strategy	48.4	70.3	80.5	56.38	(60.75	86.57	91.64	71.71
w/o K-Means	50.1	73.1	80.9	59.79	(61.04	85.52	91.34	71.69
w/o RQ-VAE	46.5	70.1	78.9	56.62	4	58.95	82.98	87.91	68.65
CART	52.6	75.4	84.2	61.77		63.58	87.91	92.84	73.59

Table 8: Ablation Study on Text-to-Audio/Video retrieval task.

Table 9: Ablation Study on MS-COCO.

Setting	R@1	R@5	R@10	MRR@10
w/o consistency loss	49.86	75.69	83.73	60.02
w/o fusion strategy	43.41	69.27	79.12	54.34
w/o K-Means	48.11	74.87	84.08	59.53
w/o RQ-VAE	51.03	76.94	84.97	62.12
CART	52.42	77.53	86.11	63.19

Table 10: Results with different number of decoder layers.

Layers	R@1	R@5	R@10	MRR@10
1	79.48	94.66	97.02	85.96
2	82.52	95.86	98.04	88.31
3	81.78	96.14	98.38	88.04
4	73.18	92.66	96.61	81.62
5	67.32	92.22	96.54	77.86

Table 11: Results for different settings of codebook dimension.

Dimensions	R@1	R@5	R@10	MRR@10
32	77.84	93.62	96.58	84.66
64	81.78	96.14	98.38	88.04
128	79.88	94.52	97.02	85.91