

IRGen: Generative Modeling for Image Retrieval

Yidan Zhang^{1,2}, Ting Zhang^{1*}, Dong Chen³, Yujing Wang³, Qi Chen³,
Xing Xie³, Hao Sun³, Weiwei Deng³, Qi Zhang³, Fan Yang³, Mao Yang³,
Qingmin Liao⁵, Jingdong Wang⁴, and Baining Guo³

¹Beijing Normal University ²The University of Tokyo ³Microsoft ⁴Baidu
⁵Shenzhen International Graduate School, Tsinghua University

Abstract. While generative modeling has become prevalent across numerous research fields, its integration into the realm of image retrieval remains largely unexplored and underjustified. In this paper, we present a novel methodology, reframing image retrieval as a variant of generative modeling and employing a sequence-to-sequence model. This approach is harmoniously aligned with the current trend towards unification in research, presenting a cohesive framework that allows for end-to-end differentiable searching. This, in turn, facilitates superior performance via direct optimization techniques. The development of our model, dubbed IRGen, addresses the critical technical challenge of converting an image into a concise sequence of semantic units, which is pivotal for enabling efficient and effective search. Extensive experiments demonstrate that our model achieves state-of-the-art performance on three widely-used image retrieval benchmarks as well as two million-scale datasets, yielding significant improvement compared to prior competitive retrieval methods. In addition, the notable surge in precision scores facilitated by generative modeling presents the potential to bypass the reranking phase, which is traditionally indispensable in practical retrieval workflows. The code is publicly available at <https://github.com/yakt00/IRGen>.

Keywords: Image Retrieval · Autoregressive Model · Generative Model

1 Introduction

Generative modeling has made significant progress in a wide range of tasks including machine translation [89], conversational modeling [1, 12, 28, 67], image captioning [108], image classification [17], text-to-image synthesis [30, 73], and many more. Originating from language and then expanding to other modalities with specially designed tokenizers, such a universal modeling approach provides a promising direction for unifying different tasks into a versatile pretrained model, which has attracted widespread attention [2, 56, 67, 94]. Yet its potential in image retrieval has been unexplored. This paper aims to take the unified trend one step further and investigates generative modeling for image retrieval.

*Corresponding author: tingzhang@bnu.edu.cn.

A practical retrieval system generally consists of two stages: feature representation learning [33, 54, 60, 105] and Approximate Nearest Neighbor (ANN) search [6, 18, 39, 47]. Most image retrieval methods focus only on one individual stage while ignoring the fact that both stages are inherently and deeply connected in actual service. Thus, the practical system often requires careful per-task hyperparameter tuning to make the most out of the coordination of the feature extraction and ANN search. While some efforts have been dedicated to joint learning of embedding and compression (or indexing), such as deep quantization / hashing [16, 19, 49, 112, 113], they still rely on a post process conducting non-exhaustive search.

In this paper, we cast image retrieval as a form of generative modeling and utilize standard Transformer architecture, as in GPT [12, 71, 72], to enable end-to-end differentiable search. Our model, IRGen, is a sequence-to-sequence model that directly generates identifiers corresponding to the given query image’s nearest neighbors. Specifically, the model takes a query image as input and autoregressively predicts discrete visual tokens, which are considered as the identifier of an image. The predicted visual tokens are supposed to point to the query image’s nearest neighbor. As such, IRGen can be trained directly from the final retrieval target starting with raw images.

Two fundamental concerns need to be addressed to enable efficient and effective image retrieval using generative modeling. First, autoregressive generative modeling is notable for its slow inference speed due to the inherently sequential nature, thus the run-time cost for retrieval increases quadratically with the length of the input sequence. Second, it is particularly difficult to model the semantic relationship between the identifiers if we drastically shortened the image identifier. Therefore, a semantic tokenizer specially designed for image retrieval is an immediate problem.

We observe that existing image tokenizers [53, 88], intended for image generation, are not suitable for image retrieval task, leading to poor performance as we will demonstrate in our experiments. We hence propose several key ingredients that first inject semantic information by applying image-level supervision rather than pixel-level reconstructive supervision, then generate dependent tokens in a sequence by leveraging the recursive property of residual quantization, and lastly ensure fast inference speed by tremendously reducing the length of the sequence via exploiting the global feature instead of spatial patch embeddings. Afterwards, we intentionally adopt the standard Transformer architecture so that we could scale up the model using existing techniques and infrastructures.

The proposed IRGen achieves state-of-the-art performance across a diverse range of image retrieval datasets, owing to its end-to-end differentiable search capability. A comprehensive empirical study over three image retrieval benchmarks proves that our model is superior to prior state-of-the-art baselines, sometimes even surpassing linear scan search methods. For instance, compared with the best baseline method, including linear scan search, our model achieves remarkable improvement, such as a 20.2% increase in precision@10 on In-shop Clothes [61], a 6.0% boost in precision@2 on CUB200 [90] and a 2.4% enhancement in preci-

sion@2 on Cars196 [50]. To assess the scalability of our model, we further experiment on million-level datasets, namely ImageNet [27] and Places365 [115], and consistently demonstrate superior performance in these challenging scenarios.

It is our belief that generative models have the potential to redefine the landscape of image retrieval. The application of generative modeling to image retrieval tasks represents an exciting opportunity to unify information retrieval across various modalities. At a technical level, our model, IRGen, effortlessly bridges the gap between feature representation learning and approximate non-exhaustive search (bypassing vector compression), creating an end-to-end differentiable framework that enables direct optimization based on ultimate retrieval goal. Furthermore, the entire framework is conceptually straightforward, with all components relying on the standard Transformer architecture, renowned for its remarkable scalability [21, 32, 79, 104].

Our contributions are concluded as:

- * To the best of our knowledge, our work represents the first effort in employing a sequence-to-sequence generative model to address image retrieval in a fully end-to-end manner that directly generates the nearest neighbor’s identifier given the query.
- * We introduce a semantic image tokenizer explicitly designed for image retrieval to obtain image identifiers, facilitating autoregressive training. It is essential to note that this tokenizer differs significantly from previous quantization models, as our identifier solely represents the image and is not employed for distance computation, as is the case in quantization.
- * Extensive experiments show that our approach exhibits substantial improvements, achieving new state-of-the-art results. We have demonstrated its scalability to million-scale datasets and provided throughput number, showcasing its near-real-time performance.

2 Related Work

Image retrieval. Traditionally, hand-crafted features are heuristically designed to describe the image content based on its color [92, 99], texture [68, 95] or shape [14]. Typical features include GIST [80], SIFT [62], SURF [8], VLAD [46] and so on. Recent years have witnessed the explosive research on deep learning based features trained over labeled images. Besides the evolvement of the network architecture designs [40, 51, 89], numerous efforts [33, 100] have been dedicated to various loss functions including classification loss [111, 117], triplet loss [110], contrastive loss [33, 48], center loss [100] and so on. The similarity is calculated through distance measure between features or evaluated through re-ranking techniques [77]. Another different line of research centers on approximate nearest neighbor search to speed up the search process, accepting a certain level of compromise in search accuracy. One way is to enable fast distance computation through hashing or quantization techniques such as LSH [41], min-Hash [22], ITQ [37], PQ [45], and many others [36, 93, 119]. The other way is to reduce the

number of distance comparison by retrieving a small number of candidates. Typical methods include partition-based indexing [6, 102] that partitions the feature space into non-overlapping clusters and graph-based indexing [44] that builds a neighborhood graph with edges connecting similar images. To improve the recall rate while ensuring fast search speed, hierarchical coarse-to-fine strategy [63] has been the popular choice that the retrieved candidates are refined level by level. Additionally, a number of works have introduced hybrid indexing [18], leveraging the best of both indexing schemes while avoiding their limitations.

Generative modeling. Deep autoregressive networks are generative sequential models that assume a product rule for factoring the joint likelihood and model each conditional distribution through a neural network. AR models have shown powerful progress in generative tasks across multiple domains such as image [17, 109], text [72, 106], audio [23, 29], and video [97, 101]. The key component involves linearizing data into a sequence of symbols with notable works such as VQ-VAE [88], RQ-VAE [53]. Recently, a number of works [25, 86, 96] further explored AR model to empower entity retrieval and document retrieval. Most related to our work are NCI [96] and DSI [86], which are concerned with document retrieval. However, these approaches utilize hierarchical k-means clustering applied to document embeddings derived from a small pretrained language model to obtain document identifiers. In contrast, we put forward a novel approach that involves learning the identifier directly from semantic supervision, and showcase its effectiveness in the context of image retrieval. We posit that this discovery can also be advantageous for document retrieval tasks.

Joint learning. There is a growing interest [19, 42, 43, 112] in the study of joint learning to improve search accuracy. For example, JPQ [113] jointly learns the deep retrieval model (equivalent to an embedding model) and the embedding index (compression). DQN [16] combines PQ and embedding learning but does not optimize the PQ clusters. DPQ [49] incorporates PQ clusters for supervised joint learning, specifically targeting image retrieval. HashGAN [31] and BGAN [84] adversarially learn binary hash codes jointly with the embedding model in an unsupervised way. However, all these prior works focus on end-to-end training of encoding and compression together while overlooking the indexing part (meaning non-exhaustive search), which is challenging to formulate as differentiable. They leverage quantization (compression index) for distance approximation but still rely on linear scan in their experiments. In contrast, our approach reframes retrieval as a completely novel end-to-end task, eliminating explicit distinctions between embedding, compression, and indexing.

3 Method

3.1 Semantic Image Tokenizer

As Transformer becomes the ubiquitous architecture in computer vision, it has emerged many successful image tokenizers such as VQ-VAE [35, 74, 88, 107], RQ-VAE [53] and so on. Despite its success in image generation, they may not be well-suited for the retrieval task for several reasons. First, the process of

decoding latent codes to reconstruct raw images is essential for generating images in synthesis tasks but is redundant for retrieval tasks. Second, the autoregressive model’s inference speed is heavily influenced by the length of the code sequence, where a shorter sequence is vital for efficient searching in our context. Presently, code sequences tend to be lengthy (for example, an 8×8 feature map with a depth of 4 in RQ-VAE leads to a sequence of 256), underscoring the need to significantly reduce sequence length. Moreover, retrieval tasks require the embedding of semantic information within the latent codes. However, the current approach to learning latent codes through image reconstruction loss primarily captures low-level details, such as inconsequential local nuances and noise, which may not be beneficial for retrieval.

Building on the insights mentioned earlier, we suggest investigating the global feature extracted from the class token, rather than relying on the default spatial tokens. This approach offers a substantial reduction in sequence length, reducing it from 64 tokens to a single token. Moreover, as a result of this approach, the class token inherently contains a condensed, high-level semantic representation. Let \mathbf{f}_{cls} denote the d -dimensional feature vector outputted from the class token, which is taken as the image representation. We adopt residual quantization (RQ) or stacked composite quantization to approximate this feature. Suppose there are M codebooks with each containing L elements, $\mathcal{C}_m = \{\mathbf{c}_{m1}, \dots, \mathbf{c}_{mL}\}$, RQ recursively maps the embedding \mathbf{f}_{cls} to a sequentially ordered M codes, $\mathbf{f}_{cls} \rightarrow \{l_1, l_2, \dots, l_M\} \in [\mathbb{L}]^M$. Let $\mathbf{r}_0 = \mathbf{f}_{cls}$, we have

$$l_m = \arg \min_{l \in [\mathbb{L}]} \|\mathbf{r}_{m-1} - \mathbf{c}_{ml}\|_2^2, \quad (1)$$

$$\mathbf{r}_m = \mathbf{r}_{m-1} - \mathbf{c}_{ml_{l_m}}, \quad m = 1, 2, \dots, M. \quad (2)$$

The process of sequentially generating discrete codes is inherently compatible with sequential autoregressive generation. This alignment helps alleviate the optimization challenges associated with modeling the relationships within identifiers. A figure illustrating the image tokenizer can be seen in the supplementary.

To further inject semantic prior, we propose to train the network under classification loss over both the original embeddings as well as the reconstructed embeddings. In particular, we consider a series of reconstruction levels denoted as $\hat{\mathbf{f}}_{cls}^{\leq m} = \sum_{i=1}^m \mathbf{c}_{il_i}$, $m = 1, 2, \dots, M$. Each prefix code thus encodes semantic information to a certain degree. Adding up the M levels of partial reconstruction error, the complete objective function is then formulated as,

$$\mathcal{L} = \mathcal{L}_{cls}(\mathbf{f}_{cls}) + \lambda_1 \sum_{m=1}^M \mathcal{L}_{cls}(\hat{\mathbf{f}}_{cls}^{\leq m}) + \lambda_2 \sum_{m=1}^M \|\mathbf{r}_m\|_2^2, \quad (3)$$

$$\mathbf{r}_m = \mathbf{f}_{cls} - \text{sg}[\hat{\mathbf{f}}_{cls}^{\leq m}], \quad m = 1, 2, \dots, M, \quad (4)$$

where \mathcal{L}_{cls} is the softmax cross-entropy loss and $\text{sg}[\cdot]$ is the stop gradient operator. During training, we adopt alternative optimization to update the codebook and the network. For computing the gradient of $\mathcal{L}_{cls}(\hat{\mathbf{f}}_{cls}^{\leq m})$, we follow the straight-through estimator [9] as in [88]. After optimization, we hope that images with

similar classes have close codes. In the experiments, we present comparison with other discrete identifiers including random codes and codes from hierarchical k-means algorithm as well as from RQ-VAE.

3.2 Encoder-Decoder for Autoregressive Retrieval

Once we have established a robust discrete latent structure equipped with semantic prior information, our next step is to train a powerful autoregressive sequence-to-sequence model solely on these discrete random variables without referring their visual content. Our encoder-decoder architecture decouples the input embedding from the generation of discrete codes. The training process involves image pairs (x_1, x_2) , where x_2 is the nearest neighbor of x_1 . Our model’s objective is to predict the identifiers of x_2 when given x_1 as input. This setup allows the model to learn the semantic relationships between images in the dataset. The left of Fig. 1 provides a concise view of our training pipeline. After training, the model begins by taking a query image as input to obtain the query embedding, which is decoded to the discrete codes. It is worth noting that the yielded discrete codes represent the query’s nearest neighbor images within the database rather than for the query.

To be specific, let the encoder be denoted as \mathbb{E} based on the ViT base architecture and the decoder be \mathbb{D} , a standard Transformer that includes causal self-attention, cross-attention and MLP layers. We leverage the spatial tokens outputted from the encoder as the input embedding, $\mathbf{e} = \mathbb{E}(x_1)$, which is injected into the decoder through cross attention. Our training objective involves predicting the next token in the image identifier sequence. Specifically, we aim to maximize the probability of the i -th token of the image identifier given the input embedding and the previously predicted tokens, $p(l_i|x_1, l_1, \dots, l_{i-1}, \theta)$, where θ denotes the parameters of both \mathbb{D} and \mathbb{E} , and l_1, l_2, \dots, l_M are the M tokens that make up the image identifier for x_2 , generated by the image tokenizer. By maximizing the probability of each token, we effectively maximize the likelihood of generating the image identifier of an image,

$$p(l_1, \dots, l_M|x_1, \theta) = \prod_{m=1}^M p(l_m|x_1, l_1, \dots, l_{m-1}, \theta). \quad (5)$$

We apply a softmax cross-entropy loss on a vocabulary of M discrete tokens during training.

3.3 Beam Search

During inference, given a query image \mathbf{q} , we first calculate the query embedding processed by the encoder \mathbb{E} and then generate discrete codes through the decoder \mathbb{D} based on the query embedding in an autoregressive manner. These discrete codes represent an image that is considered as the nearest neighbor to the query image. This illustrates the process of Top-1 retrieval by generating




Fig. 1: The left of the figure illustrates the training pipeline of encoder-decoder architecture for autoregressive retrieval. The training objective is to autoregressively predict the identifier of the query’s nearest neighbor image. The right of the figure illustrates the procedure of beam search with code length is 2 and beam size K is 3.

the sequence codes one by one according to the argmax probability. To perform Top- K retrieval, our model utilize beam search integrated into the generating process, allowing us to find the top- K images that are the closest matches.

As illustrated in right of Fig. 1, this procedure begins by defining a beam size, denoted as K in our case. Initially, the beam comprises the start-of-sequence token, and the algorithm extends the current set of candidate sequences by considering the next potential tokens in the sequence. For each candidate sequence, the algorithm outputs probabilities for multiple potential next tokens and adds them to the beam. The expanded set of sequences is then scored based on the probabilities assigned by the model, which are computed as the product of the probabilities associated with their individual elements. The top- k sequences with the highest scores are retained, while the others are discarded. This selection process, which can be implemented with a priority queue in linear time with the candidate list size, is repeated until the last token is decoded.

It’s important to note that not all generated identifiers are necessarily valid, meaning that an identifier belonging to the set $[\mathbb{L}]^M$ may not correspond to any images within the retrieval database. Therefore, we must validate the generated image identifier at each step, which can be a time-consuming process. However, we address this challenge by constructing a prefix tree containing valid codes. In the case of beam search with a prefix tree, the search process is constrained to only consider valid next tokens, as determined by the prefix tree, rather than all the possible next tokens. This constraint ensures that the search process results in valid images, largely enhancing the efficiency of the retrieval process.

Beam search vs. ANN search. Indeed, beam search can be regarded as a tree-based manner to perform approximate nearest neighbor (ANN) search. However, compared to prior methods, there is a significant difference in how they calculate the score to choose the current node. Traditional methods usually rely on calculating the distance between the query feature and the node’s associated feature, using a predefined distance metric to determine the score. Conversely, in beam search, the score or likelihood is produced by a differentiable neural network conditioned on the query. This distinction allows the retrieval process to be refined end-to-end.

Table 1: Precision comparison with different baselines, for which we consider linear scan search, Faiss IVF search and SPANN search. (repro) denotes the model reproduced by ourselves to ensure the same data process and comparable model size for fair comparison. Our model adopts beam search for retrieval, achieving significant improvement and performing even better than linear scan search.

Model	In-shop				CUB200				Cars196			
	1	10	20	30	1	2	4	8	1	2	4	8
<i>Linear scan search</i>												
Res101-Img	30.7	10.2	7.1	5.8	46.8	43.6	39.9	34.9	25.9	22.0	18.5	15.4
CLIP	57.5	22.8	16.6	14.1	66.0	63.5	59.4	53.8	70.8	67.8	63.3	57.2
CGD _(repro)	83.2	47.8	40.2	37.0	76.7	75.5	73.7	71.4	87.1	86.1	84.6	82.6
IRT _{R(repro)}	92.7	59.6	51.1	47.6	79.3	77.7	75.0	71.4	75.6	73.1	68.3	61.7
FT-CLIP	91.4	66.8	58.9	55.4	79.2	77.6	76.0	73.2	88.4	87.7	87.1	85.8
<i>Faiss IVF PQ search</i>												
CGD _(repro)	60.4	30.5	24.5	22.0	71.6	70.8	69.9	68.7	84.8	84.4	84.1	83.3
IRT _{R(repro)}	68.6	35.7	29.3	26.6	68.9	67.6	66.2	63.4	59.1	57.5	54.7	51.7
FT-CLIP	63.7	37.0	30.7	28.0	72.6	72.1	71.2	69.7	86.5	86.3	86.2	86.0
<i>ScaNN search</i>												
CGD _(repro)	83.0	47.7	40.3	37.2	76.7	75.2	73.8	71.4	87.1	86.1	84.5	82.6
IRT _{R(repro)}	92.0	58.2	50.0	46.6	79.3	77.7	75.1	71.4	75.4	72.8	68.1	61.6
FT-CLIP	90.4	64.6	56.9	53.5	79.2	77.5	76.0	73.2	88.3	87.7	87.1	85.8
<i>SPANN search</i>												
CGD _(repro)	83.0	47.7	40.3	37.1	76.7	75.5	73.7	71.4	87.0	86.1	84.6	82.6
IRT _{R(repro)}	91.4	56.2	47.9	44.5	79.3	77.6	75.0	71.4	74.8	72.4	67.6	61.1
FT-CLIP	90.2	62.9	55.1	51.8	78.5	77.6	76.0	73.2	88.6	88.1	87.5	86.3
<i>Beam search</i>												
IRGen (ours)	92.4	87.0	86.6	86.5	82.7	82.7	83.0	82.8	90.1	89.9	90.2	90.5

4 Experiments

We evaluate our method on common image retrieval datasets and on two large-scale datasets, ImageNet [27] and Places365 [115]. For a detailed description of the datasets and implementation details, please refer to the supplementary.

Baselines. We evaluate our model’s performance in comparison to five competitive baselines: 1) ResNet-101 [40] trained from ImageNet dataset, denoted as Res101-Img, which is commonly used as a feature extraction tool for various tasks; 2) CLIP [70] trained on 400M image-text pairs, known for powerful zero-shot capability; 3) CGD [48], a state-of-the-art method based on ResNet; 4) IRT [33], a Transformer-based model for image retrieval and we use the best-performing model IRT_R; 5) FT-CLIP, a baseline finetuned from CLIP on the target dataset. For both CGD and IRT, we have reproduced these models to ensure consistent data processing and comparable model sizes. Specifically, we use ResNet-101 for CGD and DeiT-B for IRT. We also provide their best results from their original papers for reference.

Search process. The baseline models primarily focus on effective feature learning. After training, these models are utilized to extract features for the database images. During the search process, a given query image is initially passed through

Table 2: Recall comparison with different baselines, for which we consider linear scan search, Faiss IVF search and SPANN search. (repro) denotes the model reproduced by ourselves to ensure the same data process and comparable model size for fair comparison. we include the best result of CGD and IRT from their original papers for context with * denotation. Our model adopt beam search for retrieval, achieving comparable performance in most cases.

Model	In-shop				CUB200				Cars196			
	1	10	20	30	1	2	4	8	1	2	4	8
Linear scan search												
Res101-Img	30.7	55.9	62.7	66.8	46.8	59.9	71.7	80.8	25.9	35.6	47	59.7
CLIP	57.5	83.0	87.5	89.7	66.0	78.1	87.7	93.5	70.8	82.6	91.1	95.9
CGD*	91.9	98.1	98.7	99.0	79.2	86.6	92.0	95.1	94.8	97.1	98.2	98.8
IRT _R *	91.9	98.1	98.7	99.0	76.6	85.0	91.1	94.3	-	-	-	-
FT-CLIP	91.4	97.3	98.1	98.5	79.2	85.0	89.3	92.0	88.4	90.5	92.5	93.8
Faiss IVF PQ search												
CGD _(repro)	60.4	76.0	77.1	77.4	71.6	77.4	81.5	84.2	84.8	88.0	89.8	91.0
IRT _{R(repro)}	68.6	79.2	80.0	80.2	68.9	77.9	85.0	89.3	59.1	70.4	78.2	83.4
FT-CLIP	63.7	70.7	71.1	71.2	72.6	78.0	82.3	85.2	86.5	86.9	87.2	87.5
ScaNN search												
CGD _(repro)	83.0	94.8	96.2	96.7	76.7	83.5	88.0	91.8	87.1	91.7	94.6	96.6
IRT _{R(repro)}	92.0	97.8	98.3	98.4	79.3	86.8	91.9	94.7	75.4	84.7	90.9	95.0
FT-CLIP	90.4	95.9	96.6	96.9	79.2	85.0	89.2	92.7	88.3	90.5	92.4	93.7
SPANN search												
CGD _(repro)	83.0	95.0	96.4	96.9	76.7	83.4	87.9	91.8	87.0	91.7	94.6	96.7
IRT _{R(repro)}	91.4	97.2	97.6	97.7	79.3	86.8	91.9	94.7	74.8	84.3	90.5	94.7
FT-CLIP	90.2	95.8	96.7	97.0	78.5	85.0	89.4	92.9	88.6	90.7	92.5	94.2
Beam search												
IRGen (ours)	92.4	96.8	97.6	97.9	82.7	86.4	89.2	91.4	90.1	92.1	93.2	93.7

the model to obtain its query feature. Subsequently, this query feature is compared with the features of database images using a specific distance metric. Following the conventions established in previous works such as [33, 48, 70], we employ the cosine distance for the CLIP model and the Euclidean distance for the other baseline models. We evaluate two search strategies: linear scan search (K-nearest neighbors or KNN) and approximate nearest neighbor search (ANN). Linear scan search is known for its accuracy but is computationally intensive. In contrast, ANN is significantly more efficient. For ANN, we explore: (i) the popular Faiss IVF PQ [47]; (ii) the state-of-the-art memory-based algorithm ScaNN [39] with the default setting; and (iii) the state-of-the-art disk-based SPANN algorithm [18]. These evaluation strategies allow us to assess the retrieval performance of our model against a variety of search methods.

4.1 Results

Tab. 1 presents a detailed performance comparison in terms of precision@ K , which assesses the percentage of retrieved candidates that share the same class as the query among the top K results. Our model consistently outperforms all

other models, even surpassing those employing linear scan search. Notably, we achieve remarkable improvements, such as a 20.2% boost in precision@10 on the In-shop Clothes dataset, a 6.0% increase in precision@2 on the CUB200 dataset, and a 2.4% gain in precision@2 on the Cars196 dataset. Furthermore, several observations can be made: 1) Finetuned models, tailored to specific datasets, exhibit significantly better performance compared to off-the-shelf feature extractors like CLIP and ImageNet-pretrained ResNet-101. 2) Generally, models equipped with ANN algorithms perform slightly worse than their counterparts using linear scan search. However, there are exceptions, such as FT-CLIP with SPANN search on the Cars196 dataset, which demonstrates the importance of end-to-end optimization. 3) Our model consistently maintains high precision scores as K increases, while other models experience a substantial drop.

Tab. 2 provides a comparison of different models using the Recall@ K metric. Recall@ K measures the proportion of queries for which at least one image among the top K retrieved candidates shares the same label as the query image, yielding a score of 1 if true and 0 otherwise. The table also includes the best recall results of CGD and IRT from their respective original papers for reference. It’s important to note that these models may have different data preprocessing, model sizes, and additional training techniques. Here are the key observations: 1) Our IRGen model achieves the highest Recall@1 score compared to all other models. However, for other recall scores, our model performs similarly or slightly worse. This discrepancy may arise from the current objective loss used in autoregressive models, which heavily optimizes for Recall@1 while giving less emphasis to other recall values. One potential solution is to incorporate the beam search process into training for joint optimization. 2) Different combinations of feature extractors and ANN algorithms exhibit significant variations across the three datasets, highlighting the challenges of achieving coordination in practical scenarios. 3) Notably, despite the high recall achieved by baselines, they often require an additional re-ranking stage to improve precision, whereas our model already attains high precision scores without the need for re-ranking.

Fig. 2 illustrates the precision-recall curve, where recall represents the true positive rate. Our approach, IRGen, consistently delivers outstanding performance, maintaining high precision and recall simultaneously. In addition to precision-recall analysis, we evaluate our model using the mean reciprocal rank (MRR) metric, which measures the inverse of the rank of the first relevant item. We compute MRR for four different values: 1, 2, 4, and 8, and display the corresponding curves in Fig. 3. The baselines employ the SPANN retrieval algorithm. Our IRGen model consistently outperforms the baselines across all evaluated metrics, confirming the effectiveness of our framework. Notably, there is significant variability in the performance gap between each baseline and our model across the three datasets, highlighting the challenges and dataset-dependent nature of retrieval tasks.

Results on million-level datasets. We further experiment our approach with ImageNet dataset [27] that contains 1,281,167 images and Places365-Standard [115] containing about 1.8M images from 365 scene categories. We compare with




Fig. 2: Precision-Recall (TPR) curve comparison for different methods on (a) In-shop Clothes, (b) CUB200 and (c) Cars196 dataset.




Fig. 3: MRR with respect to 1,2,4,8 comparison for different methods on (a) In-shop Clothes, (b) CUB200 and (c) Cars196 dataset.

the strong baselines including CLIP model as well as FT-CLIP model finetuned based on CLIP model. The comparison is reported in Fig. 4 and Tab. 3, focusing on precision@ K and MAP@100. Our IRGen model consistently outperforms the baselines, achieving the best results in terms of precision@ K and MAP@100. The precision values for our model remain consistently high as K increases, while the baselines experience noticeable performance degradation. These results confirm the effectiveness of our model in handling large-scale datasets like ImageNet, where it maintains high precision across varying values of K and outperforms the baseline models.

Comparison to supervised deep quantization methods. We make the comparison with recent supervised deep quantization methods on ImageNet, including the state-of-the-art methods ADSVQ [116] and DPQ [49]. Notably, these methods exclusively utilize 100 classes from ImageNet, as indicated in their original papers, making it a substantially easier case. According to the original paper, DPQ [49] employed ResNet-50 (25.56M parameters), ADSVQ [116] utilized AlexNet (60.3M parameters), and DTQ [58] also used AlexNet. To ensure fair comparison, we employ a smaller network, ViT-S (22.1M parameters) and rerun over 100 classes. Tab. 4 illustrates the superiority of our approach.

4.2 Ablations

The effect of identifiers. In our study of image identifiers, we compared four different approaches: (1) assigning random identifiers to images, (2) hierarchical k-means (HKM), (3) Product Quantization (PQ) and (4) using the image tokenizer RQ-VAE [11]. The results of this comparison are summarized in Tab. 6.

Table 3: MAP@100 on two million-level datasets. The results of CLIP and FT-CLIP are retrieved by SPANN. **Table 4:** MAP@1000 for 32 bits (equivalent to the length of our identifier) comparison on ImageNet dataset.

Model	Dataset		Method	MAP@1000 for 32 bits
	ImageNet	Places365		
CLIP	44.1	22.1	DTQ [58]	68.12
FT-CLIP	65.5	30.3	ADSVQ [116]	75.21
IRGen(Ours)	76.0	44.3	DPQ [49]	87.70
			IRGen(Ours)	93.52

Table 5: Generalize to new data. We split 5% of the training data from the ImageNet dataset for inference and remained unseen during training.

Model	Precision		
	1	10	100
FT-CLIP + Linear Scan	70.6	65.0	55.6
IRGen (Ours)	77.0	77.9	77.4

The random assignment of identifiers to images yielded expectedly lower performance. This performance gap can be attributed to the fact that models with random identifiers need to learn not only the interaction between queries and image identifiers but also allocate capacity to learn relationships within the identifiers themselves. On the other hand, HKM showed superior performance compared to random assignment, underscoring the significance of semantic identifiers. However, our proposed semantic image identifiers demonstrated a clear improvement over HKM, highlighting their effectiveness in enhancing retrieval performance. PQ codes perform worse than the random assigned identifiers. This underperformance may stem from the model’s challenges in constructing semantic information from PQ codes, which might be incompatible with autoregressive prediction mechanism. This highlights the significance of our sequence order in autoregressive training process. In contrast, the performance of RQ-VAE significantly trailed behind our model, with a performance less than 10 percent. We attribute this difference to the sequence length in RQ-VAE which is too long for the model to effectively capture relationships within the identifiers.

Generalize to new data. Addressing the inclusion of fresh data holds particular significance, especially in the context of search scenarios. To assess this capacity, we conduct an experiment where we intentionally withheld 5% of the training data from the ImageNet dataset during the training phase and introduced it during inference, all without updating the existing codebook and AR model. In this experiment, we compare our model with the formidable baseline FT-CLIP, which is equipped with a linear scan search. The results, as displayed in Tab. 5, reveal that our model maintains superior performance even when confronted with new data. This observation highlights our model’s ability to effectively generalize to previously unseen data. In scenarios where more new data come in with vastly different semantic structure, leading to the distribution of gallery data change drastically, all data-dependent component such as embedding model as well as joint learning methods including IRGen, have to




Fig. 4: Precision comparison on large scale datasets: ImageNet and Places365.

Table 6: Ablation study on the image identifier (T=length).

Identifier	T	Precision				Recall			
		1	10	20	30	1	10	20	30
Random	4	87.6	75.4	70.8	68.3	87.6	95.1	96.0	96.1
HKM ₁₀₀	4	88.2	80.0	78.2	77.3	87.2	93.1	94.3	95.0
HKM ₂₀₀	3	89.0	81.6	79.8	79.0	89.0	93.9	94.9	95.8
HKM ₅₀₀	2	89.5	81.7	79.8	78.9	89.5	95.3	96.5	97.0
PQ	4	73.8	66.1	61.8	59.2	73.8	90.2	93.3	94.5
Ours	4	92.1	89.7	89.5	89.5	92.1	96.8	97.5	97.7


re-train the model and re-construct the data index. Essentially, generalization (or dynamic updates) only supports a small number of new data that do not substantially alter the distribution which we have demonstrated in our case. In future we may extend IRGen to an adaptive structure capable of capturing the evolving distribution of new data during training.

Storage requirement and Inference throuput. We first provide a comprehensive comparison of model capacity on the ImageNet dataset and the comparison is shown in the Tab. 7. It is crucial to emphasize that a retrieval system’s storage considerations extend beyond just the model; it must also store features for all database vectors, constituting a substantial portion of storage requirements, especially in the context of a large database. While compression can alleviate this concern, the performance is severely compromised due to quantization. Our GPU inference cost, tested on a Tesla P100 PCIe 16GB GPU, is 0.037s/image, which is near-real-time, considering practical retrieval times at the order of 10 milliseconds. Moreover, our approach achieves 76% MAP@100, Moreover, our approach achieves 76% MAP@100, significantly surpassing the FT-CLIP with SPANN result of 65.5%, and is competitive with FT-CLIP utilizing linear scan, which achieves 77%. Notably, however, the FT-CLIP with linear scan method demands 0.14 seconds per image and considerable memory usage.

We further report additional inference throughput on Inshop dataset for 100 queries, with beam sizes set at 1, 10, 20, and 30 for comparison. This analysis is conducted on an NVIDIA V100-16G GPU and we break down the time cost of each component during autoregressive retrieval. As shown in Fig. 5, the encoder

Table 7: Comprehensive comparison of storage requirement and inference cost.

Method	Model(Params)	Database storage
IRGen	ViT-base encoder(86M)+24-layer decoder(307M)	4.9MB: 4 bytes/image
FT-CLIP	ViT-base(86M)	3.67GB: 768x4 bytes/image

**Fig. 5:** Illustrating the search speed using beam search.

is quite fast, while the autoregressive decoder becomes the major bottleneck, especially as the beam size increases. Additional time is consumed for checking the validity of predictions, as it’s possible that the predicted identifier may not exist in the database. Overall, the time cost is within an acceptable range. For instance, it takes approximately 0.07 seconds (with a beam size of 10) or 0.19 seconds (with a beam size of 30) per query.

For inference speed, we anticipate that the prevalent reliance on GPUs will swiftly lead to newer GPU generations with increased processing power, more efficient memory management, and improved hardware support, enhancing the speed of our autoregressive retrieval. Additionally, our model, being GPT-based, can leverage rapidly developed acceleration techniques designed for GPT.

5 Conclusion

In this paper, we delve into the realm of generative modeling to enable end-to-end image retrieval, a process that directly connects a query image to its closest match. With the introduction of semantic image tokenizer, we’ve demonstrated that our model excels at achieving remarkable precision without compromising recall. Through extensive ablation studies and evaluations, we’ve underscored the superior performance of our approach. We believe that this innovative approach to generative modeling in image retrieval not only pushes the boundaries of this field but also holds potential for broader applications.

Limitations. While our model has shown significant performance improvements, it’s important to acknowledge its limitations. First, though our model has demonstrated scalability to million-scale datasets, dealing with billion-scale datasets is a complex challenge. It may necessitate even larger models with higher capacity, potentially impacting inference speed. Striking a balance between model capacity and speed is an area that warrants exploration for efficient and effective billion-scale search. Second Training large autoregressive models requires substantial computational resources, which raises environmental concerns. Research efforts to enable efficient training, such as fast fine-tuning of pretrained models, are crucial to mitigate energy consumption and environmental impact.

Acknowledgements

This work is supported by Fundamental Research Funds for the Central Universities (2243100004).

References

1. Adiwardana, D., Luong, M.T., So, D.R., Hall, J., Fiedel, N., Thoppilan, R., Yang, Z., Kulshreshtha, A., Nemade, G., Lu, Y., et al.: Towards a human-like open-domain chatbot. arXiv preprint arXiv:2001.09977 (2020)
2. Alayrac, J.B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., et al.: Flamingo: a visual language model for few-shot learning. arXiv preprint arXiv:2204.14198 (2022)
3. Alzu'bi, A., Amira, A., Ramzan, N.: Semantic content-based image retrieval: A comprehensive study. *Journal of Visual Communication and Image Representation* **32**, 20–54 (2015)
4. Andoni, A., Indyk, P.: Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Communications of the ACM* **51**(1), 117–122 (2008)
5. Babenko, A., Lempitsky, V.: Additive quantization for extreme vector compression. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 931–938 (2014)
6. Babenko, A., Lempitsky, V.: The inverted multi-index. *IEEE transactions on pattern analysis and machine intelligence* **37**(6), 1247–1260 (2014)
7. Bai, S., Tang, P., Torr, P.H., Latecki, L.J.: Re-ranking via metric fusion for object retrieval and person re-identification. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 740–749 (2019)
8. Bay, H., Tuytelaars, T., Van Gool, L.: Surf: Speeded up robust features. In: *European conference on computer vision*. pp. 404–417. Springer (2006)
9. Bengio, Y., Léonard, N., Courville, A.: Estimating or propagating gradients through stochastic neurons for conditional computation. arXiv preprint arXiv:1308.3432 (2013)
10. Bentley, J.L.: K-d trees for semidynamic point sets. In: *Proceedings of the sixth annual symposium on Computational geometry*. pp. 187–197 (1990)
11. Bevilacqua, M., Ottaviano, G., Lewis, P., Yih, W.t., Riedel, S., Petroni, F.: Autoregressive search engines: Generating substrings as document identifiers. arXiv preprint arXiv:2204.10628 (2022)
12. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in neural information processing systems* **33**, 1877–1901 (2020)
13. Cao, B., Araujo, A., Sim, J.: Unifying deep local and global features for image search. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*. pp. 726–743. Springer (2020)
14. Cao, Y., Wang, C., Zhang, L., Zhang, L.: Edgel index for large-scale sketch-based image search. In: *CVPR 2011*. pp. 761–768. IEEE (2011)
15. Cao, Y., Long, M., Wang, J., Liu, S.: Collective deep quantization for efficient cross-modal retrieval. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 31 (2017)

16. Cao, Y., Long, M., Wang, J., Zhu, H., Wen, Q.: Deep quantization network for efficient image retrieval. In: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence. p. 3457–3463. AAAI’16, AAAI Press (2016)
17. Chen, M., Radford, A., Child, R., Wu, J., Jun, H., Luan, D., Sutskever, I.: Generative pretraining from pixels. In: International conference on machine learning. pp. 1691–1703. PMLR (2020)
18. Chen, Q., Zhao, B., Wang, H., Li, M., Liu, C., Li, Z., Yang, M., Wang, J.: Spann: Highly-efficient billion-scale approximate nearest neighborhood search. *Advances in Neural Information Processing Systems* **34**, 5199–5212 (2021)
19. Chen, T., Li, L., Sun, Y.: Differentiable product quantization for end-to-end embedding compression. In: International Conference on Machine Learning. pp. 1617–1626. PMLR (2020)
20. Cho, J., Lu, J., Schwenk, D., Hajishirzi, H., Kembhavi, A.: X-lxmert: Paint, caption and answer questions with multi-modal transformers. arXiv preprint arXiv:2009.11278 (2020)
21. Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H.W., Sutton, C., Gehrmann, S., et al.: Palm: Scaling language modeling with pathways. arXiv preprint arXiv:2204.02311 (2022)
22. Chum, O., Philbin, J., Zisserman, A., et al.: Near duplicate image detection: Min-hash and tf-idf weighting. In: *Bmvc*. vol. 810, pp. 812–815 (2008)
23. Chung, Y.A., Hsu, W.N., Tang, H., Glass, J.: An unsupervised autoregressive model for speech representation learning. arXiv preprint arXiv:1904.03240 (2019)
24. De Cao, N., Aziz, W., Titov, I.: Highly parallel autoregressive entity linking with discriminative correction. arXiv preprint arXiv:2109.03792 (2021)
25. De Cao, N., Izacard, G., Riedel, S., Petroni, F.: Autoregressive entity retrieval. arXiv preprint arXiv:2010.00904 (2020)
26. De Cao, N., Wu, L., Popat, K., Artetxe, M., Goyal, N., Plekhanov, M., Zettlemoyer, L., Cancedda, N., Riedel, S., Petroni, F.: Multilingual autoregressive entity linking. *Transactions of the Association for Computational Linguistics* **10**, 274–290 (2022)
27. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
28. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
29. Dhariwal, P., Jun, H., Payne, C., Kim, J.W., Radford, A., Sutskever, I.: Jukebox: A generative model for music. arXiv preprint arXiv:2005.00341 (2020)
30. Ding, M., Zheng, W., Hong, W., Tang, J.: Cogview2: Faster and better text-to-image generation via hierarchical transformers. arXiv preprint arXiv:2204.14217 (2022)
31. Dizaji, K.G., Zheng, F., Sadoughi, N., Yang, Y., Deng, C., Huang, H.: Unsupervised deep generative adversarial hashing network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3664–3673 (2018)
32. Du, N., Huang, Y., Dai, A.M., Tong, S., Lepikhin, D., Xu, Y., Krikun, M., Zhou, Y., Yu, A.W., Firat, O., et al.: Glam: Efficient scaling of language models with mixture-of-experts. In: International Conference on Machine Learning. pp. 5547–5569. PMLR (2022)
33. El-Nouby, A., Neverova, N., Laptev, I., Jégou, H.: Training vision transformers for image retrieval. arXiv preprint arXiv:2102.05644 (2021)

34. Erin Liang, V., Lu, J., Wang, G., Moulin, P., Zhou, J.: Deep hashing for compact binary codes learning. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2475–2483 (2015)
35. Gafni, O., Polyak, A., Ashual, O., Sheynin, S., Parikh, D., Taigman, Y.: Make-a-scene: Scene-based text-to-image generation with human priors. arXiv preprint arXiv:2203.13131 (2022)
36. Ge, T., He, K., Ke, Q., Sun, J.: Optimized product quantization. IEEE transactions on pattern analysis and machine intelligence **36**(4), 744–755 (2013)
37. Gong, Y., Lazebnik, S., Gordo, A., Perronnin, F.: Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. IEEE transactions on pattern analysis and machine intelligence **35**(12), 2916–2929 (2012)
38. Gulrajani, I., Kumar, K., Ahmed, F., Taiga, A.A., Visin, F., Vazquez, D., Courville, A.: Pixelvae: A latent variable model for natural images. arXiv preprint arXiv:1611.05013 (2016)
39. Guo, R., Sun, P., Lindgren, E., Geng, Q., Simcha, D., Chern, F., Kumar, S.: Accelerating large-scale inference with anisotropic vector quantization. In: International Conference on Machine Learning. pp. 3887–3896. PMLR (2020)
40. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
41. Indyk, P., Motwani, R.: Approximate nearest neighbors: towards removing the curse of dimensionality. In: Proceedings of the thirtieth annual ACM symposium on Theory of computing. pp. 604–613 (1998)
42. Jang, Y.K., Cho, N.I.: Generalized product quantization network for semi-supervised image retrieval. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3420–3429 (2020)
43. Jang, Y.K., Cho, N.I.: Self-supervised product quantization for deep unsupervised image retrieval. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12085–12094 (2021)
44. Jayaram Subramanya, S., Devvrit, F., Simhadri, H.V., Krishnawamy, R., Kadekodi, R.: Diskann: Fast accurate billion-point nearest neighbor search on a single node. Advances in Neural Information Processing Systems **32** (2019)
45. Jegou, H., Douze, M., Schmid, C.: Product quantization for nearest neighbor search. IEEE transactions on pattern analysis and machine intelligence **33**(1), 117–128 (2010)
46. Jégou, H., Douze, M., Schmid, C., Pérez, P.: Aggregating local descriptors into a compact image representation. In: 2010 IEEE computer society conference on computer vision and pattern recognition. pp. 3304–3311. IEEE (2010)
47. Johnson, J., Douze, M., Jégou, H.: Billion-scale similarity search with gpus. IEEE Transactions on Big Data **7**(3), 535–547 (2019)
48. Jun, H., Ko, B., Kim, Y., Kim, I., Kim, J.: Combination of multiple global descriptors for image retrieval. arXiv preprint arXiv:1903.10663 (2019)
49. Klein, B., Wolf, L.: End-to-end supervised product quantization for image search and retrieval. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5036–5045 (2019). <https://doi.org/10.1109/CVPR.2019.00518>
50. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: Proceedings of the IEEE international conference on computer vision workshops. pp. 554–561 (2013)

51. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *Communications of the ACM* **60**(6), 84–90 (2017)
52. Lagunes-Fortiz, M., Damen, D., Mayol-Cuevas, W.: Centroids triplet network and temporally-consistent embeddings for in-situ object recognition. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 10796–10802. IEEE (2020)
53. Lee, D., Kim, C., Kim, S., Cho, M., Han, W.S.: Autoregressive image generation using residual quantization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11523–11532 (2022)
54. Lee, S., Seong, H., Lee, S., Kim, E.: Correlation verification for image retrieval. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5374–5384 (2022)
55. Lew, M.S., Sebe, N., Djeraba, C., Jain, R.: Content-based multimedia information retrieval: State of the art and challenges. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* **2**(1), 1–19 (2006)
56. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. arXiv preprint arXiv:2301.12597 (2023)
57. Li, X., Uricchio, T., Ballan, L., Bertini, M., Snoek, C.G., Bimbo, A.D.: Socializing the semantic gap: A comparative survey on image tag assignment, refinement, and retrieval. *ACM Computing Surveys (CSUR)* **49**(1), 1–39 (2016)
58. Liu, B., Cao, Y., Long, M., Wang, J., Wang, J.: Deep triplet quantization. In: Proceedings of the 26th ACM international conference on Multimedia. pp. 755–763 (2018)
59. Liu, Y., Zhang, D., Lu, G., Ma, W.Y.: A survey of content-based image retrieval with high-level semantics. *Pattern recognition* **40**(1), 262–282 (2007)
60. Liu, Z., Cheng, K.T., Huang, D., Xing, E.P., Shen, Z.: Nonuniform-to-uniform quantization: Towards accurate quantization via generalized straight-through estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4942–4952 (2022)
61. Liu, Z., Luo, P., Qiu, S., Wang, X., Tang, X.: Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1096–1104 (2016)
62. Lowe, D.G.: Object recognition from local scale-invariant features. In: Proceedings of the seventh IEEE international conference on computer vision. vol. 2, pp. 1150–1157. Ieee (1999)
63. Malkov, Y.A., Yashunin, D.A.: Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and machine intelligence* **42**(4), 824–836 (2018)
64. Martinez, J., Hoos, H.H., Little, J.J.: Stacked quantizers for compositional vector compression. arXiv preprint arXiv:1411.2173 (2014)
65. Noh, H., Araujo, A., Sim, J., Weyand, T., Han, B.: Large-scale image retrieval with attentive deep local features. In: Proceedings of the IEEE international conference on computer vision. pp. 3456–3465 (2017)
66. Norouzi, M., Fleet, D.J.: Cartesian k-means. In: Proceedings of the IEEE Conference on computer Vision and Pattern Recognition. pp. 3017–3024 (2013)
67. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C.L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al.: Training language models to follow instructions with human feedback. arXiv preprint arXiv:2203.02155 (2022)

68. Park, M., Jin, J.S., Wilson, L.S.: Fast content-based image retrieval using quasi-gabor filter and reduction of image feature dimension. In: Proceedings fifth IEEE southwest symposium on image analysis and interpretation. pp. 178–182. IEEE (2002)
69. Qiu, Z., Yao, T., Mei, T.: Deep quantization: Encoding convolutional activations with deep generative model. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6759–6768 (2017)
70. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. pp. 8748–8763. PMLR (2021)
71. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al.: Improving language understanding by generative pre-training (2018)
72. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. OpenAI blog **1**(8), 9 (2019)
73. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125 (2022)
74. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: International Conference on Machine Learning. pp. 8821–8831. PMLR (2021)
75. Razavian, A.S., Sullivan, J., Carlsson, S., Maki, A.: Visual instance retrieval with deep convolutional networks. *ITE Transactions on Media Technology and Applications* **4**(3), 251–258 (2016)
76. Ren, J., Zhang, M., Li, D.: Hm-ann: Efficient billion-point nearest neighbor search on heterogeneous memory. *Advances in Neural Information Processing Systems* **33**, 10672–10684 (2020)
77. Revaud, J., Almazán, J., Rezende, R.S., Souza, C.R.d.: Learning with average precision: Training image retrieval with a listwise loss. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5107–5116 (2019)
78. Sharif Razavian, A., Azizpour, H., Sullivan, J., Carlsson, S.: Cnn features off-the-shelf: an astounding baseline for recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 806–813 (2014)
79. Shoyebi, M., Patwary, M., Puri, R., LeGresley, P., Casper, J., Catanzaro, B.: Megatron-lm: Training multi-billion parameter language models using model parallelism. arXiv preprint arXiv:1909.08053 (2019)
80. Siagian, C., Itti, L.: Rapid biologically-inspired scene classification using features shared with visual attention. *IEEE transactions on pattern analysis and machine intelligence* **29**(2), 300–312 (2007)
81. Siméoni, O., Avrithis, Y., Chum, O.: Local features and visual words emerge in activations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11651–11660 (2019)
82. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
83. Smeulders, A.W., Worring, M., Santini, S., Gupta, A., Jain, R.: Content-based image retrieval at the end of the early years. *IEEE Transactions on pattern analysis and machine intelligence* **22**(12), 1349–1380 (2000)
84. Song, J., He, T., Gao, L., Xu, X., Hanjalic, A., Shen, H.T.: Binary generative adversarial networks for image retrieval. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)

85. Tan, F., Yuan, J., Ordonez, V.: Instance-level image retrieval using reranking transformers. In: proceedings of the IEEE/CVF international conference on computer vision. pp. 12105–12115 (2021)
86. Tay, Y., Tran, V.Q., Dehghani, M., Ni, J., Bahri, D., Mehta, H., Qin, Z., Hui, K., Zhao, Z., Gupta, J., et al.: Transformer memory as a differentiable search index. arXiv preprint arXiv:2202.06991 (2022)
87. Teichmann, M., Araujo, A., Zhu, M., Sim, J.: Detect-to-retrieve: Efficient regional aggregation for image search. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5109–5118 (2019)
88. Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. *Advances in neural information processing systems* **30** (2017)
89. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
90. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset (2011)
91. Wang, J., Wang, J., Song, J., Xu, X.S., Shen, H.T., Li, S.: Optimized cartesian k-means. *IEEE Transactions on Knowledge and Data Engineering* **27**(1), 180–192 (2014)
92. Wang, J., Hua, X.S.: Interactive image search by color map. *ACM Transactions on Intelligent Systems and Technology (TIST)* **3**(1), 1–23 (2011)
93. Wang, J., Zhang, T.: Composite quantization. *IEEE transactions on pattern analysis and machine intelligence* **41**(6), 1308–1322 (2018)
94. Wang, W., Bao, H., Dong, L., Bjorck, J., Peng, Z., Liu, Q., Aggarwal, K., Mohammed, O.K., Singhal, S., Som, S., et al.: Image as a foreign language: Beit pre-training for all vision and vision-language tasks. arXiv preprint arXiv:2208.10442 (2022)
95. Wang, X.Y., Zhang, B.B., Yang, H.Y.: Content-based image retrieval by integrating color and texture features. *Multimedia tools and applications* **68**(3), 545–569 (2014)
96. Wang, Y., Hou, Y., Wang, H., Miao, Z., Wu, S., Sun, H., Chen, Q., Xia, Y., Chi, C., Zhao, G., et al.: A neural corpus indexer for document retrieval. arXiv preprint arXiv:2206.02743 (2022)
97. Weissenborn, D., Täckström, O., Uszkoreit, J.: Scaling autoregressive video models. arXiv preprint arXiv:1906.02634 (2019)
98. Wen, Y., Zhang, K., Li, Z., Qiao, Y.: A discriminative feature learning approach for deep face recognition. In: European conference on computer vision. pp. 499–515. Springer (2016)
99. Wengert, C., Douze, M., Jégou, H.: Bag-of-colors for improved image search. In: Proceedings of the 19th ACM international conference on Multimedia. pp. 1437–1440 (2011)
100. Wiczorek, M., Michalowski, A., Wroblewska, A., Dabrowski, J.: A strong baseline for fashion retrieval with person re-identification models. In: International Conference on Neural Information Processing. pp. 294–301. Springer (2020)
101. Wu, C., Liang, J., Ji, L., Yang, F., Fang, Y., Jiang, D., Duan, N.: Nüwa: Visual synthesis pre-training for neural visual world creation. In: European Conference on Computer Vision. pp. 720–736. Springer (2022)
102. Xia, Y., He, K., Wen, F., Sun, J.: Joint inverted indexing. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3416–3423 (2013)

103. Xiao, T., Li, H., Ouyang, W., Wang, X.: Learning deep feature representations with domain guided dropout for person re-identification. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1249–1258 (2016)
104. Xu, Y., Lee, H., Chen, D., Hechtman, B., Huang, Y., Joshi, R., Krikun, M., Lepikhin, D., Ly, A., Maggioni, M., et al.: Gspmd: general and scalable parallelization for ml computation graphs. arXiv preprint arXiv:2105.04663 (2021)
105. Yang, M., He, D., Fan, M., Shi, B., Xue, X., Li, F., Ding, E., Huang, J.: Dolg: Single-stage image retrieval with deep orthogonal fusion of local and global features. In: Proceedings of the IEEE/CVF International conference on Computer Vision. pp. 11772–11781 (2021)
106. Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R.R., Le, Q.V.: Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems* **32** (2019)
107. Yu, J., Li, X., Koh, J.Y., Zhang, H., Pang, R., Qin, J., Ku, A., Xu, Y., Baldrige, J., Wu, Y.: Vector-quantized image modeling with improved vqgan. arXiv preprint arXiv:2110.04627 (2021)
108. Yu, J., Wang, Z., Vasudevan, V., Yeung, L., Seyedhosseini, M., Wu, Y.: Coca: Contrastive captioners are image-text foundation models. arXiv preprint arXiv:2205.01917 (2022)
109. Yu, J., Xu, Y., Koh, J.Y., Luong, T., Baid, G., Wang, Z., Vasudevan, V., Ku, A., Yang, Y., Ayan, B.K., et al.: Scaling autoregressive models for content-rich text-to-image generation. arXiv preprint arXiv:2206.10789 (2022)
110. Yuan, Y., Chen, W., Yang, Y., Wang, Z.: In defense of the triplet loss again: Learning robust person re-identification with fast approximated triplet loss and label distillation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 354–355 (2020)
111. Zhai, A., Wu, H.Y.: Classification is a strong baseline for deep metric learning. arXiv preprint arXiv:1811.12649 (2018)
112. Zhan, J., Mao, J., Liu, Y., Guo, J., Zhang, M., Ma, S.: Jointly optimizing query encoder and product quantization to improve retrieval performance. In: Proceedings of the 30th ACM International Conference on Information & Knowledge Management. pp. 2487–2496 (2021)
113. Zhang, H., Shen, H., Qiu, Y., Jiang, Y., Wang, S., Xu, S., Xiao, Y., Long, B., Yang, W.Y.: Joint learning of deep retrieval model and product quantization based embedding index. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 1718–1722 (2021)
114. Zhang, L., Rui, Y.: Image search—from thousands to billions in 20 years. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* **9**(1s), 1–20 (2013)
115. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2017)
116. Zhou, C., Po, L.M., Ou, W.: Angular deep supervised vector quantization for image retrieval. *IEEE Transactions on Neural Networks and Learning Systems* **33**(4), 1638–1649 (2020)
117. Zhou, K., Yang, Y., Cavallaro, A., Xiang, T.: Omni-scale feature learning for person re-identification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3702–3712 (2019)

118. Zhou, W., Li, H., Tian, Q.: Recent advance in content-based image retrieval: A literature survey. arXiv preprint arXiv:1706.06064 (2017)
119. Zhu, H., Long, M., Wang, J., Cao, Y.: Deep hashing network for efficient similarity retrieval. In: Proceedings of the AAAI conference on Artificial Intelligence. vol. 30 (2016)