

# SEEING THROUGH WORDS: CONTROLLING VISUAL RETRIEVAL QUALITY WITH LANGUAGE MODELS

Jianglin Lu<sup>1,2\*</sup> Simon Jenni<sup>1</sup> Kushal Kafle<sup>1</sup> Jing Shi<sup>1</sup> Handong Zhao<sup>1</sup> Yun Fu<sup>2</sup>

<sup>1</sup>Adobe Research <sup>2</sup>Northeastern University

## ABSTRACT

Text-to-image retrieval is a fundamental task in vision-language learning, yet in real-world scenarios it is often challenged by short and underspecified user queries. Such queries are typically only one or two words long, rendering them semantically ambiguous, prone to collisions across diverse visual interpretations, and lacking explicit control over the quality of retrieved images. To address these issues, we propose a new paradigm of *quality-controllable retrieval*, which enriches short queries with contextual details while incorporating explicit notions of image quality. Our key idea is to leverage a generative language model as a query completion function, extending underspecified queries into descriptive forms that capture fine-grained visual attributes such as pose, scene, and aesthetics. We introduce a general framework that conditions query completion on discretized quality levels, derived from relevance and aesthetic scoring models, so that query enrichment is not only semantically meaningful but also quality-aware. The resulting system provides three key advantages: ① *flexibility*, it is compatible with any pretrained vision-language model (VLMs) without modification; ② *transparency*, enriched queries are explicitly interpretable by users; and ③ *controllability*, enabling retrieval results to be steered toward user-preferred quality levels. Extensive experiments demonstrate that our proposed approach significantly improves retrieval results and provides effective quality control, *bridging the gap between the expressive capacity of modern VLMs and the underspecified nature of short user queries*. Our code is available at <https://github.com/Jianglin954/QCQC>.

## 1 INTRODUCTION

Text-to-image retrieval (T2IR) aims to return the most relevant images from a gallery given a textual query. Recent progress in this task has been largely driven by vision–language models (VLMs) (Jia et al., 2021; Yang et al., 2022; Li et al., 2023; Yang et al., 2025; Lu et al., 2025a; Huang et al., 2025), which learn joint representations of text and images through large-scale pretraining on web-scale image–text pairs (Schuhmann et al., 2021; 2022; Liu et al., 2023). These models significantly narrow the semantic gap between modalities and achieve strong alignment across diverse benchmarks (Ilharco et al., 2021; Singh et al., 2022; Li et al., 2024; Lu et al., 2025b; Dong et al., 2026).

Despite these advances, retrieval performance often degrades in realistic scenarios where user queries are very short (typically just one or two words, e.g., “a dog”). Short queries encode only limited semantics, which results in large and ambiguous search subspaces and less discriminative results. This issue becomes more pronounced in large-scale galleries, where underspecified queries yield many candidate matches and cause semantic collisions among visually diverse results.

Another limitation of existing retrieval systems is their singular focus on semantic alignment. Naïve retrieval approaches simply return the top- $k$  images with the highest similarity scores, overlooking other critical aspects of user satisfaction such as aesthetics, interestingness, or popularity (Yi et al., 2023; Abdullahi & Grabner, 2024; Wang et al., 2025). In practice, *retrieval quality* is context-dependent: art students may prefer visually inspiring images, architects may seek unique and creative references, and shoppers may favor popular or visually appealing products. However, conventional retrieval systems lack mechanisms for steering retrieval toward these quality dimensions.

\*Corresponding author: jianglinlu@outlook.com. Work done during an internship at Adobe Research.

To address these limitations, we introduce the task of *quality-controllable retrieval* (QCR). Formally, given a frozen VLM and a short textual query, the objective is to retrieve images that not only align semantically but also satisfy user-specified quality requirements. This setting is feasible because short queries occupy a broad region of the embedding space that contains images of varying perceptual quality. With appropriate conditioning, this region can be partitioned into perceptually distinct subsets, enabling fine-grained quality-aware retrieval.

In this work, we define retrieval quality along two widely applicable dimensions: *relevance* (semantic consistency) (Cherti et al., 2023) and *aesthetics* (visual appeal) (Yi et al., 2023). For each image in the gallery, we construct auxiliary annotations consisting of a textual description, a relevance score, and an aesthetic score. We discretize these continuous scores into categorical quality levels and associate each description with its corresponding quality condition.

The central challenge is how to steer retrieval results toward specific quality levels given short queries. We propose a simple yet effective solution: *quality-conditioned query completion* (QCQC). QCQC enriches short queries with quality-aware details by leveraging a generative large language model (LLM). Trained on the quality-augmented dataset, the LLM learns to append appropriate descriptive phrases that capture both semantic and quality-related attributes. By conditioning on distinct quality levels, QCQC generates targeted query completions that steer retrieval toward specific regions of the embedding space. This capability is particularly valuable in practice, as users often struggle to articulate quality preferences or may lack a clear understanding of what constitutes “high” or “low” quality within a dataset. By modeling how textual descriptions vary across quality levels, our approach bridges this gap and enables more controllable, quality-aware retrieval through conditioned query completion. Our key contributions in this work can be summarized as follows:

- A new problem: we introduce quality-controllable retrieval, a new setting where retrieval can be explicitly conditioned on user-defined quality requirements.
- A general solution: we propose QCQC, a generative query completion framework that leverages LLMs to enrich short queries with quality-aware descriptive details.
- Validation: we conduct extensive experiments to show that QCQC effectively steers retrieval outcomes according to quality preferences and is compatible to multiple VLMs.

## 2 PRELIMINARIES

### 2.1 MOTIVATION

We study the problem of text-to-image retrieval, where the goal is to return the desired images from a large gallery given a set of natural language queries. Specifically, let  $\mathcal{Q} := \{Q_1, \dots, Q_m\}$  denote a collection of  $m$  text queries and  $\mathcal{I} := \{I_1, \dots, I_n\}$  an image gallery of size  $n$ . We consider a state-of-the-art VLM as the retrieval backbone, equipped with a text encoder  $g : \mathcal{Q} \rightarrow \mathbb{R}^d$  and an image encoder  $f : \mathcal{I} \rightarrow \mathbb{R}^d$ , both producing  $d$ -dimensional normalized embeddings. Given a query set  $\mathcal{Q}$ , the system returns the top- $\eta$  relevant images according to

$$\mathcal{X} := \text{sort}(f(\mathcal{I}), g(\mathcal{Q}), \eta), \quad (1)$$

where  $\mathcal{X} \subseteq \mathcal{I}$  denotes the top- $\eta$  matches of queries  $\mathcal{Q}$ . The sort function typically operates on the similarity scores  $\mathbf{S} \in \mathbb{R}^{m \times n}$  with  $S_{ij} := g(Q_i)^\top f(I_j)$ .

Although modern VLMs achieve strong cross-modal alignment, retrieval performance deteriorates in realistic scenarios where user queries are usually very short (typically just one or two words, e.g., “a dog”). Given such short queries, naïve retrieval system faces several challenges: ① *Semantic ambiguity*: a few words can refer to a wide range of possible images, leading to a large and diffuse search subspace with less discriminative retrieval results. ② *Semantic collisions*: short queries tend to yield close similarity scores for visually diverse images. These collisions confuse ranking and are particularly problematic in large-scale galleries where many candidate images match the vague query. ③ *Lack of quality control*: the quality of retrieved images is not explicitly enforced during retrieval. At best, one can apply post-retrieval filtering, but the system itself provides no mechanism to ensure that high-quality results consistently appear among the top matches. These issues highlight *a fundamental gap between the expressive capacity of modern VLMs and the underspecified nature of user queries*, motivating the need for query enrichment and controllable retrieval mechanisms.

## 2.2 PROBLEM SETTING

To address the above limitations, we propose to enrich short queries with additional descriptive details that potentially capture more distinguishable attributes of images. Formally, let  $h$  denote a query completion function that maps  $Q$  to enriched queries  $h(Q)$ . Retrieval is then performed as

$$\tilde{\mathcal{X}} := \text{sort}(f(\mathcal{I}), g(h(Q)), \eta), \quad (2)$$

where  $h(Q)$  augments the short queries with contextual details. The enriched queries are expected to capture not only object categories but also additional information such as pose, scene, action, and fine-grained attributes. To be effective, the completion function *should be aware of the retrieval gallery, so that it generates meaningful context rather than irrelevant content.*

To achieve this, we implement  $h$  using a generative large language model (LLM). However, simply training the LLM on image descriptions is insufficient, since it cannot guarantee that retrieval results satisfy user expectations of quality. Instead, we partition the textual descriptions into non-overlapped quality levels  $\mathcal{C}$  that reflect different image quality categories. We then finetune the LLM with these quality levels, enabling it to generate query completions conditioned on quality preferences. This yields the formulation of our *quality-controllable retrieval (QCR)* :

$$\tilde{\mathcal{X}} := \text{sort}(f(\mathcal{I}), g(\text{LLM}(Q; \mathcal{C})), \eta), \quad (3)$$

where  $\text{LLM}(Q; \mathcal{C})$  expands the short queries based on the specified quality constraint  $\mathcal{C}$ . The extended queries thus steer retrieval toward images that align with the desired quality criteria.

This approach offers several practical benefits: ① *Flexibility*: it requires no modification to pre-trained VLMs and remains compatible with any VLMs; ② *Transparency*: the generated query completions are human-readable, allowing users to review and select preferred options. ③ *Controllability*: the LLM can produce different query completions according to distinct quality conditions  $\mathcal{C}$ , enabling explicit quality control during retrieval. In the following section, we provide theoretical justification for why enriching short queries may improve retrieval performance.

## 2.3 THEORETICAL ANALYSIS

We model query completion as a structured perturbation and analyze its effect on the similarity matrices  $\mathbf{S}$  through the lens of rank variation under perturbations.

Let  $Q^+ = \{Q_1^+, \dots, Q_m^+\} := h(Q)$  denote the extended queries by  $h$ , where  $Q_i^+ := Q_i + \text{suffix}_i$ ,  $\forall i \in \{1, \dots, m\}$ , and  $\text{suffix}_i$  denotes additional descriptive details appended to query  $Q_i$ . Let  $\mathbf{C} \in \mathbb{R}^{n \times d}$  be the image embedding matrix with rows  $\mathbf{c}_j := f(I_j) \in \mathbb{R}^d$ ,  $\forall j \in \{1, \dots, n\}$ , and  $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{m \times d}$  be two sets of text embeddings with a strict one-to-one pairing of rows, with rows  $\mathbf{a}_i := g(Q_i) \in \mathbb{R}^d$  and  $\mathbf{b}_i := g(Q_i^+) \in \mathbb{R}^d$ ,  $\forall i \in \{1, \dots, m\}$ . Let  $r := \text{rank}(\mathbf{A})$  be the rank of  $\mathbf{A}$ ,  $\sigma_r(\mathbf{A})$  be the smallest nonzero singular value of  $\mathbf{A}$ , and  $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$  denote its singular value decomposition (SVD). We then partition the right singular vectors as  $\mathbf{V} = [\mathbf{V}_S \ \mathbf{V}_\perp]$ , where  $\mathbf{V}_S \in \mathbb{R}^{d \times r}$  and  $\mathbf{V}_\perp \in \mathbb{R}^{d \times (d-r)}$  satisfy  $\text{span}(\mathbf{V}_S) = \mathcal{R}(\mathbf{A})$  and  $\text{span}(\mathbf{V}_\perp) = \mathcal{R}(\mathbf{A})^\perp$ , with  $\mathcal{R}(\mathbf{A}) := \text{span}\{\mathbf{a}_1^\top, \dots, \mathbf{a}_m^\top\} \subseteq \mathbb{R}^d$  the row space of  $\mathbf{A}$ .

**Definition 1.** We define a perturbation matrix  $\mathbf{\Delta} := \mathbf{B} - \mathbf{A} \in \mathbb{R}^{m \times d}$ , score matrices  $\mathbf{S}_A := \mathbf{A}\mathbf{C}^\top \in \mathbb{R}^{m \times n}$  and  $\mathbf{S}_B := \mathbf{B}\mathbf{C}^\top \in \mathbb{R}^{m \times n}$  for the queries  $Q$  and  $Q^+$ ,  $\mathbf{A}_S := \mathbf{A}\mathbf{V}_S$ ,  $\mathbf{\Delta}_S := \mathbf{\Delta}\mathbf{V}_S$ ,  $\mathbf{\Delta}_\perp := \mathbf{\Delta}\mathbf{V}_\perp$ ,  $\mathbf{C}_S := \mathbf{C}\mathbf{V}_S$ ,  $\mathbf{C}_\perp := \mathbf{C}\mathbf{V}_\perp$ ,  $\mathbf{X} := (\mathbf{A}_S + \mathbf{\Delta}_S)\mathbf{C}_S^\top$ ,  $\mathbf{Y} := \mathbf{\Delta}_\perp\mathbf{C}_\perp^\top$ ,  $\mathcal{U} := \text{col}(\mathbf{X})$ , and  $\mathbf{P} := \mathbf{P}_X$  as the orthogonal projector onto  $\mathcal{U}$ .

**Lemma 1.** If  $\text{rank}(\mathbf{X}_I) = r$  and  $\|\mathbf{X}_I^\dagger \mathbf{P} \mathbf{Y}_I\|_2 < 1$ , then  $\text{rank}(\mathbf{X}_I + \mathbf{P}\mathbf{Y}_I) = r$ .

**Proposition 1.** Assume that: i)  $\|\mathbf{\Delta}\|_2 < \sigma_r(\mathbf{A})$ ; ii) there exists  $I \subseteq \{1, \dots, n\}$  with  $|I| = r$  such that the columns  $\mathbf{X}_I$  form a basis of  $\mathcal{U}$ ; iii)  $\|\mathbf{X}_I^\dagger \mathbf{P} \mathbf{Y}_I\|_2 < 1$ ; and iv) there exists disjoint index set  $K \subseteq \{1, \dots, n\} \setminus I$  such that  $k := \text{rank}\left((\mathbf{I} - \mathbf{P}_{Z_I})\mathbf{Z}_K\right) \geq 1$ , where  $\mathbf{Z} := (\mathbf{I} - \mathbf{P})\mathbf{Y}$ ,  $\mathbf{Z}_I := \mathbf{Z}_{:,I}$ ,  $\mathbf{Z}_K := \mathbf{Z}_{:,K}$ , and  $\mathbf{P}_{Z_I}$  be the orthogonal projector onto  $\text{col}(\mathbf{Z}_I)$ . Then,  $\text{rank}(\mathbf{S}_B) > \text{rank}(\mathbf{S}_A)$ .

This proposition indicates that, under certain conditions, the score matrix  $\mathbf{S}_B$  derived from query completion can express more independent scoring patterns and has the ability to potentially make finer-grained distinctions. The detailed proof and analysis are provided in the Appendix A.1.




Figure 1: Aesthetic and relevance score distributions of Flickr2.4M in (a) and (b), and of MS-COCO in (c) and (d). It is worth noting that the numbers of high-quality and low-quality images are limited, which leads to the average scores of any two random sets being very close.

### 3 QUALITY-CONDITIONED QUERY COMPLETION

#### 3.1 QUALITY DEFINITION

For the proposed QCR task, we require a clear notion of *quality*. In this work, we characterize quality along two primary dimensions: ① *Relevance*, which measures the semantic consistency between textual queries and their corresponding images; and ② *Aesthetics*, which reflects the visual appeal or attractiveness of retrieved images. Note that the proposed QCQC framework is inherently flexible, permitting the incorporation of arbitrary quality metric, provided that reliable scoring models are available and applicable to general image datasets. Other notions of quality, such as *interestingness* (Gygli et al., 2013; Abdullahu & Grabner, 2024) can also be adopted in a similar manner and are left for future exploration. To facilitate user control over retrieved results, we discretize the quality dimensions into non-overlapping conditions. Specifically, we define  $\mathcal{C}^R$  for relevance and  $\mathcal{C}^A$  for aesthetics, each partitioned into perceptually distinct and user-friendly levels. For example, both can be represented as  $\mathcal{C}^R, \mathcal{C}^A := \{\text{Low}, \text{Medium}, \text{High}\}$ .

#### 3.2 DATA GENERATION

To ensure that the completion function can perceive the retrieval gallery, we construct an augmented training dataset for each gallery  $\mathcal{I}$ . The dataset integrates three complementary components: textual descriptions  $\mathcal{T} = \{T_i\}_{i=1}^n$  of images, relevance scores  $\mathbf{s}^r \in \mathbb{R}^n$ , and aesthetic scores  $\mathbf{s}^a \in \mathbb{R}^n$ .

**Textual Descriptions.** For each image  $I_i$ , we generate a textual description  $T_i$  using an image caption model  $\text{CAP}(\cdot)$ , i.e.,  $T_i = \text{CAP}(I_i), \forall i \in \{1, \dots, n\}$ . In our experiments, we utilize strong pretrained captioning models without additional fine-tuning for description generation. Each  $T_i$  is a concise sentence summarizing the main content of the image.

**Aesthetic Scores.** We assign an aesthetic score  $s_i^a$  to each image  $I_i$  using an aesthetic evaluation model  $\text{EV}_A(\cdot)$ , i.e.,  $s_i^a = \text{EV}_A(I_i), \forall i \in \{1, \dots, n\}$ . The aesthetic scores represent the visual quality of the images, with higher scores indicating greater visual appeal.

**Relevance Scores.** For each image-description pair  $\{I_i, T_i\}$ , we compute a relevance score using a pretrained VLM. Specifically, we extract the image feature  $f(I_i)$  and text feature  $g(T_i)$ , then calculate their cosine similarity as their relevance score, i.e.,  $s_i^r = \cos(f(I_i), g(T_i)), \forall i \in \{1, \dots, n\}$ .

#### 3.3 TRAINING FRAMEWORK

**Score Discretization.** To simulate quality-controlled retrieval, we discretize the continuous quality scores of images into categorical levels that are more intuitive for users. Given a score vector  $\mathbf{r}$  (either aesthetics  $\mathbf{s}^a$  or relevance  $\mathbf{s}^r$ ), each score  $r_i$  is mapped into one of three descriptive levels by partitioning the score distribution into three percentiles:<sup>1</sup>

$$l(\mathbf{r}_i) = \begin{cases} \text{Low}, & r_i \leq \text{perc}(\mathbf{r}, p_1), \\ \text{High}, & r_i > \text{perc}(\mathbf{r}, p_2), \\ \text{Medium}, & \text{otherwise.} \end{cases} \quad (4)$$

<sup>1</sup>Our framework is general and supports arbitrary numbers of levels depending on the desired granularity. In Table 5, Sec. 4.5, we provide an example with five quality levels.

Table 1: Query completions with their retrieved images and quality scores on MS-COCO

Rel: <b>Low</b> , Aes: <b>Low</b>	Rel: <b>Medium</b> , Aes: <b>Medium</b>	Rel: <b>High</b> , Aes: <b>High</b>
 <i>a train that is sitting on the tracks in gravel</i> Aes 4.715, Rel 0.347	 <i>a train sitting on the tracks with black smoke coming out of it</i> Aes 4.818, Rel 0.382	 <i>a train is traveling near some water and houses</i> Aes 5.935, Rel 0.394
 <i>a bird standing on the ground near some leaves</i> Aes 4.616, Rel 0.346	 <i>a bird flying above some brown water</i> Aes 5.079, Rel 0.374	 <i>a bird flying across some water at the beach</i> Aes 5.120, Rel 0.386
 <i>a teddy bear wearing eye glasses and laying on a bed</i> Aes 4.788, Rel 0.359	 <i>a teddy bear that is sitting on a tree</i> Aes 5.649, Rel 0.388	 <i>a teddy bear sitting on a wall next to an old stone house</i> Aes 5.818, Rel 0.437

Here,  $r_i$  is the score of the  $i$ -th sample and  $\text{perc}(r, p)$  calculates the  $p\%$  percentile of  $r$  as

$$\text{perc}(r, p) = \tilde{r}[\lfloor \xi \rfloor] + (\xi - \lfloor \xi \rfloor) \cdot (\tilde{r}[\lfloor \xi \rfloor + 1] - \tilde{r}[\lfloor \xi \rfloor]), \quad (5)$$

where  $\xi = \frac{p}{100} \cdot (n - 1)$ ,  $\tilde{r}$  is the sorted version of  $r$ , and  $\lfloor \cdot \rfloor$  is the floor function. Figure 1 illustrates the distributions of aesthetics and relevance scores and their discretized partitions.




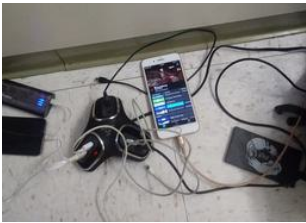

**Instruction Design.** We then train the completion function LLM on the augmented training set  $\mathcal{D} = \{\mathcal{T}, s^a, s^r\}$ . The discretized quality levels serve as explicit conditions within instructions, enabling the LLM to generate quality-aware query completions. For each image  $I_i$ , we design a concise instruction  $P_i$  of the following form:

"Relevance:  $l(s_i^r)$ , Aesthetic:  $l(s_i^a)$ , Query: "

where  $l(s_i^r)$  and  $l(s_i^a)$  represent the categorical quality levels defined in Eq. (4). During training, this instruction provides a lightweight yet effective mechanism to condition query generation on the specified quality preferences.

**Model Training.** To simulate the quality control process during model training, we use the descriptive levels  $l(s_i^r)$  and  $l(s_i^a)$  of image  $I_i$  as the quality conditions, which are incorporated into the instruction  $P_i$ . We then concatenate instructions  $P_i$  with the textual description  $T_i$  for each image  $I_i$ , and then train the proposed QCQC model with the standard autoregressive next-token predic-

Table 2: Query completions with their retrieved images and quality scores on Flickr2.4M

Aes: <b>Low</b> , Rel: <b>Low</b>	Aes: <b>Medium</b> , Rel: <b>Medium</b>	Aes: <b>High</b> , Rel: <b>High</b>
 <i>a chair with wires on it</i> Aes 4.019, Rel 0.362	 <i>a chair with red and black ropes on it</i> Aes 4.847, Rel 0.379	 <i>a chair on a stage in a field</i> Aes 5.257, Rel 0.387
 <i>a cell phone with wires attached to it</i> Aes 4.531, Rel 0.362	 <i>a cell phone with an acoustic guitar on it</i> Aes 5.035, Rel 0.390	 <i>a cell phone on a tripod in front of a waterfall in yellowstone national park</i> Aes 5.441, Rel 0.429
 <i>a bowl of soup with meat and vegetables in it</i> Aes 4.648, Rel 0.343	 <i>a bowl on display</i> Aes 4.980, Rel 0.379	 <i>a bowl with flowers on it</i> Aes 5.386, Rel 0.387

tion loss. In this way, our QCQC model learns to generate query completions that are not only semantically relevant but also controllable according to the given quality constraints.

**Inference Strategy.** During the inference stage, we concatenate a similar instruction with each testing query. To simulate user preferences, we evaluate various relevance-aesthetic combinations, such as “low relevance, low aesthetic” and “high relevance, high aesthetic”. Then the model generates completed queries based on the instructions, testing queries, and specified quality conditions. For efficient similarity search on large-scale galleries, we utilize the FAISS library (Johnson et al., 2019) to identify the nearest images for the queries.

## 4 EXPERIMENTS

### 4.1 EXPERIMENTAL SETTINGS

**Datasets.** We evaluate our method on two image datasets: one with real textual descriptions and one without. For the image-only one, we construct a large dataset sourced from the Openverse website (Openverse, 2025). We refer to this dataset as Flickr2.4M, which contains over 2.4 million CC0-licensed images randomly selected from the Flickr subset of Openverse. For image

Table 3: Retrieval quality of various methods on Flickr2.4M. CoCa and Blip2 are used to generate textual descriptions; **L** (Low), **M** (Medium), and **H** (High) indicate the quality conditions; and Ctrl specifies whether the method enables controllable retrieval over quality. For both average relevance (Ave Rel) and average aesthetics (Ave Aes), higher values indicate better retrieval quality.

Quality	VLM	Aes Cond Rel Cond	<b>L</b> <b>L</b>	<b>L</b> <b>M</b>	<b>L</b> <b>H</b>	<b>M</b> <b>L</b>	<b>M</b> <b>M</b>	<b>M</b> <b>H</b>	<b>H</b> <b>L</b>	<b>H</b> <b>M</b>	<b>H</b> <b>H</b>	Ctrl ?
Prefix	--	Ave Aes Ave Rel	4.735 0.350	4.735 0.350	4.735 0.350	4.735 0.350	4.735 0.350	4.735 0.350	4.735 0.350	4.735 0.350	4.735 0.350	×
LLaMA3	--	Ave Aes Ave Rel	4.730 0.351	4.822 0.351	4.831 0.351	4.823 0.353	4.837 0.351	4.784 0.350	4.798 0.354	4.722 0.354	4.842 0.352	×
GPT-4o	--	Ave Aes Ave Rel	4.359 0.378	4.651 0.361	4.728 0.357	4.712 0.358	4.668 0.360	4.791 0.356	4.791 0.361	4.816 0.357	5.056 0.361	×
PT	--	Ave Aes Ave Rel	4.776 0.345	4.556 0.346	4.722 0.349	4.811 0.349	4.781 0.346	4.757 0.348	4.693 0.345	4.751 0.350	4.746 0.350	×
FT	CoCa	Ave Aes Ave Rel	4.756 0.365	4.834 0.368	4.777 0.364	4.838 0.363	4.863 0.369	4.882 0.368	4.821 0.365	4.905 0.364	4.770 0.365	×
Ours	CoCa	Ave Aes Ave Rel	<b>4.458</b> <b>0.355</b>	<b>4.615</b> <b>0.366</b>	<b>4.530</b> <b>0.391</b>	<b>4.934</b> <b>0.354</b>	<b>4.852</b> <b>0.360</b>	<b>4.841</b> <b>0.386</b>	<b>5.222</b> <b>0.353</b>	<b>5.170</b> <b>0.368</b>	<b>5.270</b> <b>0.390</b>	✓
FT	Blip2	Ave Aes Ave Rel	4.795 0.370	4.871 0.368	4.890 0.367	4.894 0.367	4.844 0.371	4.856 0.366	4.901 0.367	4.847 0.371	4.888 0.369	×
Ours	Blip2	Ave Aes Ave Rel	<b>4.541</b> <b>0.353</b>	<b>4.523</b> <b>0.370</b>	<b>4.455</b> <b>0.397</b>	<b>4.940</b> <b>0.354</b>	<b>4.906</b> <b>0.366</b>	<b>4.922</b> <b>0.396</b>	<b>5.309</b> <b>0.355</b>	<b>5.222</b> <b>0.372</b>	<b>5.191</b> <b>0.390</b>	✓

datasets with real textual descriptions, we adopt the widely-used MS-COCO dataset for experiments, which includes both images and human-annotated descriptions. Specifically, we utilize the training subset of MS-COCO, which consists of 118, 287 samples, each sample containing one image and five corresponding descriptions. In total, approximately 0.6 million descriptions are used for training.

**Model Selection.** For the backbone of our method, we evaluate two different LLMs: GPT2-1.5B (Radford et al., 2019) and Qwen2.5-0.5B (Yang et al., 2025). Other LLMs can be validated similarly and we leave them for future study. We adopt two caption models CAP( $\cdot$ ), including a pretrained CoCa (Yu et al., 2022) and a pretrained Blip2 (Li et al., 2023) model. For feature extraction, we adopt a pretrained VLM OpenCLIP (ViT-H-14-quickgelu) (Cherti et al., 2023; Ilharco et al., 2021). The relevance score is computed as the cosine similarity between the features of each image-description pair. For  $EV_A(\cdot)$ , we use a pretrained aesthetic predictor (Schuhmann, 2022).

**Implementation Details.** For GPT2-1.5B (Radford et al., 2019), we set the learning rate, warmup steps, number of epochs, and batch size to  $2e - 3$ , 100, 50, 150, respectively. For Qwen2.5-0.5B (Yang et al., 2025), these hyperparameters are set to  $2e - 5$ , 100, 30, and 80, respectively. For score discretization, we set  $p_1 = 33$  and  $p_2 = 66$  to divide the score distribution into three evenly spaced percentiles (examples of five-level cases are also considered). Note that we only train the query completion model LLM, while the quality evaluation model  $EV_A(\cdot)$ , the caption models CAP( $\cdot$ ), and the retrieval model VLM are all pretrained without additional fine-tuning. Since the pretrained caption models may occasionally generate non-English characters, we clean these characters directly before training to prevent potential issues for query completion. Before training, we prepend a start token to the instructions and append an end token to the descriptions. The training loss is computed only on the tokens of the descriptions and the end tokens, while excluding those of the instructions.

**Evaluation Strategy.** For performance evaluation, we use the 80 class names from MS-COCO dataset as query objectives. These include common objects such as trains, cars, and animals, as well as more specific categories like teddy bear, fire hydrant, and toothbrush. Based on the capitalization of each class name, we prepend either "a" or "an" to form the input queries. Since we focus on controlling the quality of retrieved images, we use aesthetic and relevance scores as the evaluation metrics. We calculate and report the average aesthetic and relevance scores of the retrieved images across all input queries as the final evaluation performance.

Table 4: Retrieval quality of various methods on MS-COCO, where **L** (Low), **M** (Medium), and **H** (High) indicate the quality conditions for retrieval, and Ctrl specifies whether the method enables controllable retrieval over image quality. For both average relevance (Ave Rel) and average aesthetics (Ave Aes), higher values indicate better retrieval quality.

Quality	Aes Cond Rel Cond	<b>L</b>	<b>L</b>	<b>L</b>	<b>M</b>	<b>M</b>	<b>M</b>	<b>H</b>	<b>H</b>	<b>H</b>	Ctrl ?
		<b>L</b>	<b>M</b>	<b>H</b>	<b>L</b>	<b>M</b>	<b>H</b>	<b>L</b>	<b>M</b>	<b>H</b>	
Prefix	Ave Aes	4.817	4.817	4.817	4.817	4.817	4.817	4.817	4.817	4.817	×
	Ave Rel	0.349	0.349	0.349	0.349	0.349	0.349	0.349	0.349	0.349	
LLaMA3	Ave Aes	4.903	4.891	4.855	4.916	4.875	4.880	4.871	4.858	4.911	×
	Ave Rel	0.348	0.349	0.347	0.348	0.349	0.347	0.348	0.350	0.344	
GPT-4o	Ave Aes	4.673	4.754	4.686	4.782	4.808	4.880	4.838	5.075	5.048	×
	Ave Rel	0.371	0.357	0.354	0.360	0.358	0.350	0.359	0.352	0.351	
PT	Ave Aes	4.819	4.793	4.789	4.829	4.810	4.828	4.794	4.826	4.820	×
	Ave Rel	0.343	0.340	0.344	0.348	0.339	0.344	0.346	0.343	0.340	
FT	Ave Aes	4.925	4.845	4.848	4.882	4.934	4.990	4.849	4.941	4.929	×
	Ave Rel	0.370	0.367	0.366	0.368	0.368	0.365	0.371	0.371	0.367	
FT-CoCa	Ave Aes	4.878	4.852	4.859	4.902	4.858	4.941	4.952	4.961	4.944	×
	Ave Rel	0.346	0.351	0.356	0.349	0.350	0.354	0.345	0.352	0.352	
FT-Blip2	Ave Aes	4.828	4.815	4.785	4.932	4.894	4.893	5.034	4.948	4.933	×
	Ave Rel	0.350	0.352	0.356	0.344	0.351	0.353	0.345	0.351	0.347	
Ours	Ave Aes	<b>4.811</b>	<b>4.790</b>	<b>4.773</b>	<b>4.911</b>	<b>4.873</b>	<b>4.862</b>	<b>5.016</b>	<b>4.983</b>	<b>5.024</b>	✓
	Ave Rel	<b>0.356</b>	<b>0.370</b>	<b>0.382</b>	<b>0.354</b>	<b>0.370</b>	<b>0.387</b>	<b>0.352</b>	<b>0.365</b>	<b>0.387</b>	

## 4.2 QUALITATIVE VALIDATION

We first perform qualitative analysis to validate whether our approach effectively achieves quality control in retrieval. In Tables 1 and 2, we present three retrieved images per query, along with their corresponding completed queries and quality scores under three different quality conditions, including “low relevance, low aesthetic”, “medium relevance, medium aesthetic”, and “high relevance, high aesthetic”. As shown, our method generates distinct query completions for different quality conditions. From left to right, as the quality level improves, both aesthetic and relevance scores increase accordingly. This demonstrates that our proposed method effectively controls the quality of the retrieved images. We provide more qualitative results with failure cases in the Appendix A.5.

## 4.3 QUANTITATIVE VALIDATION

Since there are no existing text-to-image retrieval methods that can be directly applicable to the proposed QCR task, we design the following baselines for quantitative comparison: a) *Prefix*: using original short queries directly without query completion; b) *PT (Pretrained)*: using a pretrained LLM for query completion without finetuning; c) *FT (Finetuned)*: finetuning a pretrained LLM on textual descriptions while conditioning on randomly generated quality scores; and d) *general-purpose LLMs*, including the LLaMA-3 (LLaMA-3-8B-Instruct) (Grattafiori et al., 2024) and GPT-4o (via official API) (OpenAI et al., 2024).

Tables 3 and 4 report the retrieval performance of the baseline models and our proposed method with Qwen2.5 (Yang et al., 2025) on the two datasets, respectively. From these tables, we make the following key observations. ① Prefix-only retrieval yields unsatisfactory quality performance, highlighting the necessity of query completion. ② Pretrained models (without finetuning) for query completion degrade retrieval quality, performing worse than using only the query prefix in most cases. This is because these pretrained models tend to generate irrelevant words, negatively impacting retrieval performance. ③ Finetuning on textual descriptions improves both relevance and aesthetics compared to prefix-only and pretrained models. However, models finetuned on randomly assigned scores fail to effectively control the quality of the retrieved images. ④ Our method not only enhances retrieval under high-quality conditions but also excels in quality control, demonstrating strong adaptability regardless of whether it is trained on real or generated captions.

Table 5: Results with five quality levels.

$\mathcal{M}$		Relevance (Red $\rightarrow$ Red)				
		VL	L	M	H	VH
Aesthetics (Green) $\leftarrow$ Green	VL	4.597 0.355	4.507 0.364	4.610 0.375	4.529 0.382	4.445 0.397
	L	4.805 0.353	4.765 0.366	4.825 0.369	4.729 0.380	4.761 0.392
Aesthetics (Green)	M	4.909 0.355	4.961 0.3642	4.878 0.370	4.889 0.3754	4.901 0.390
	H	5.028 0.355	4.967 0.365	5.045 0.370	4.952 0.374	5.009 0.387
Aesthetics (Green)	VH	5.282 0.355	5.153 0.363	5.263 0.371	5.245 0.378	5.121 0.389

Table 6: Comparison with post-retrieval filtering, where the *rerank* method first retrieves the top- $k$  images based on relevance and then reorders the candidates by aesthetic scores to identify the best result.

		$k$	1	2	3	5	10
Rerank	Aes	4.735	4.947	5.014	5.198	5.313	
	Rel	0.350	0.348	0.347	0.345	0.341	
LLaMA3	Aes	4.842	5.071	5.154	5.298	5.377	
	Rel	0.352	0.349	0.347	0.342	0.337	
GPT-4o	Aes	5.056	5.205	5.293	5.393	5.518	
	Rel	0.361	0.356	0.353	0.349	0.343	
Ours	Aes	5.236	5.320	5.364	5.432	5.533	
	Rel	0.387	0.385	0.381	0.376	0.366	

#### 4.4 DATASET DEPENDENCE

To achieve quality control in retrieval, the model should be tailored to the specific dataset, as different datasets exhibit varying quality characteristics. To illustrate this, we conduct cross-dataset retrieval experiments. Specifically, we evaluate retrieval quality on MS-COCO using queries completed by the model finetuned on Flickr2.4M. In Table 4, we assess FT-CoCa and FT-Blip2, which are finetuned on descriptions generated by CoCa and Blip2, respectively. The results indicate that both models achieve higher aesthetic scores as quality conditions improve, suggesting that aesthetically relevant semantic cues may be universal across natural images. Nevertheless, they consistently exhibit low relevance across all quality conditions. This limitation stems from the dataset mismatch between the query completion and image retrieval stages, since the two datasets encode different semantic information. See Appendix A.3 for additional analysis and results.

#### 4.5 FURTHER VALIDATION

Table 5 presents the results on the Flickr2.4M dataset across five quality levels: VL (Very-Low), L (Low), M (Medium), H (High), and VH (Very-High). As shown, our method effectively enables fine-grained control over the quality of retrieved images, adhering to more nuanced descriptive constraints. We also compare against a post-retrieval filtering baseline that first retrieves images based on relevance and then re-ranks the results by aesthetic scores. The comparison results are listed Table 6. As shown, this two-stage strategy is unreliable for short queries, which typically offer vague representations and limited descriptive cues. As a result, the initial retrieval set tends to be semantically broad and aesthetically subpar, leaving little room for the re-ranking step to improve. While increasing  $k$  can surface images with higher aesthetic quality, it typically comes at the cost of reduced semantic relevance, illustrating a trade-off between these two quality dimensions. In contrast, our method performs quality control during the query stage, which inherently guides retrieval toward the desired quality level. This quality-aware conditioning cannot be achieved by the two-step baseline, which lacks knowledge of the dataset’s quality distribution and operates in a detached, post-hoc manner. See Appendix A.5 for more experimental results.

## 5 CONCLUSION

We presented a quality-controllable retrieval framework to address the limitations of short and under-specified text queries in text-to-image retrieval. Our key idea is to enrich queries using a generative language model conditioned on discretized quality levels, enabling retrieval that is both semantically expressive and aligned with user quality preferences. Extensive experiments demonstrate that our proposed QCQC approach effectively improves retrieval quality, serving as a flexible augmentation to existing VLMs while enabling quality control in retrieval. Future work will extend our method to other dimensions of quality beyond relevance and aesthetics, such as interestingness, diversity, or user personalization. We hope this work inspires further research on integrating controllable language-based query enrichment with large-scale multimodal retrieval systems.

## STATEMENTS

### ETHICS STATEMENT

This work investigates large language models for query completion in text-to-image retrieval, where image quality information is integrated into the training process. The study relies on publicly available datasets and does not involve human subjects, private information, or sensitive content. We acknowledge that retrieval models may inherit biases present in the underlying vision–language datasets; however, our approach does not introduce new data and instead focuses on methodological contributions. The models and results are intended solely for academic research, and no harmful or deceptive applications are pursued. We adhere to the ICLR Code of Ethics and confirm that this research complies with principles of fairness, transparency, and responsible use.

### REPRODUCIBILITY STATEMENT

We are committed to ensuring the reproducibility of our work. We have released our code at GitHub. Theoretical results are stated with all necessary assumptions in the main text and the complete proofs are provided in the Appendix. Experimental settings are included in the main body. Together, these resources are intended to enable full replication and verification of our results.

### REFERENCES

- Fitim Abdullahu and Helmut Grabner. Commonly interesting images. *arXiv preprint arXiv:2409.16736*, 2024.
- Ziv Bar-Yossef and Naama Kraus. Context-sensitive query auto-completion. In *Proceedings of the 20th International Conference on World Wide Web*, pp. 107–116. Association for Computing Machinery, 2011.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, pp. 1877–1901, 2020.
- Fei Cai and Maarten De Rijke. A survey of query auto completion in information retrieval. *Foundations and Trends in Information Retrieval*, 10(4):273–363, 2016.
- Yudong Chen, Sen Wang, Jianglin Lu, Zhi Chen, Zheng Zhang, and Zi Huang. Local graph convolutional networks for cross-modal hashing. In *Proceedings of the 29th ACM international conference on multimedia*, pp. 1921–1928, 2021.
- Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for contrastive language-image learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2818–2829, 2023.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1*, pp. 4171–4186. Association for Computational Linguistics, 2019.
- Qihua Dong, Luis Figueroa, Handong Zhao, Kushal Kafle, Jason Kuen, Zhihong Ding, Scott Cohen, and Yun Fu. Cot referring: Improving referring expression tasks with grounded reasoning. *arXiv preprint arXiv:2510.06243*, 2025.
- Qihua Dong, Kuo Yang, Lin Ju, Handong Zhao, Yitian Zhang, Yizhou Wang, Huimin Zeng, Jianglin Lu, and Yun Fu. Ref-adv: Exploring MLLM visual reasoning in referring expression tasks. In *The Fourteenth International Conference on Learning Representations*, 2026.

- Yuting Gao, Jinfeng Liu, Zihan Xu, Jun Zhang, Ke Li, and Chunhua Shen. Pyramidclip: hierarchical feature alignment for vision-language model pretraining. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, 2022.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, and Ahmad Al-Dahle et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Michael Gygli, Helmut Grabner, Hayko Riemenschneider, Fabian Nater, and Luc Van Gool. The interestingness of images. In *Proceedings of the IEEE international conference on computer vision*, pp. 1633–1640, 2013.
- Yiyang Huang, Yizhou Wang, and Yun Fu. D-CoDe: Scaling image-pretrained VLMs to video via dynamic compression and question decomposition. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 11798–11811. Association for Computational Linguistics, November 2025.
- Yiyang Huang, Liang Shi, Yitian Zhang, Yi Xu, and Yun Fu. SHIELD: Suppressing hallucinations in LVLM encoders via bias and vulnerability defense. In *The Fourteenth International Conference on Learning Representations*, 2026.
- Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishal Shankar, Hongseok Namkoong, John Miller, Hannaneh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Openclip. [https://github.com/mlfoundations/open\\_clip](https://github.com/mlfoundations/open_clip), 2021. GitHub repository.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pp. 4904–4916. PMLR, 2021.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3):535–547, 2019.
- Dong-Ho Lee, Zhiqiang Hu, and Roy Ka-Wei Lee. Improving text auto-completion with next phrase prediction. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pp. 4434–4438. Association for Computational Linguistics, 2021.
- Yibin Lei, Yu Cao, Tianyi Zhou, Tao Shen, and Andrew Yates. Corpus-steered query expansion with large language models. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics*, pp. 393–401. Association for Computational Linguistics, 2024.
- Jinhao Li, Haopeng Li, Sarah Monazam Erfani, Lei Feng, James Bailey, and Feng Liu. Visual-text cross alignment: Refining the similarity score in vision-language models. In *Forty-first International Conference on Machine Learning*, 2024.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pp. 19730–19742. PMLR, 2023.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *Advances in Neural Information Processing Systems*, volume 36, pp. 34892–34916, 2023.
- Jianglin Lu, Hailing Wang, Jie Zhou, Mengfan Yan, and Jiajun Wen. Deep asymmetric hashing with dual semantic regression and class structure quantization. *Information Sciences*, 589:235–249, 2022.
- Jianglin Lu, Hailing Wang, Yi Xu, Yizhou Wang, Kuo Yang, and Yun Fu. Representation potentials of foundation models for multimodal alignment: A survey. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 16669–16684. Association for Computational Linguistics, 2025a.

- Jianglin Lu, Hailing Wang, Kuo Yang, Yitian Zhang, Simon Jenni, and Yun Fu. The indra representation hypothesis for multimodal alignment. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025b.
- Mayug Maniparambil, Raiymbek Akshulakov, Yasser Abdelaziz Dahou Djilali, Mohamed El Amine Seddik, Sanath Narayan, Karttikeya Mangalam, and Noel E. O'Connor. Do vision and language encoders represent the world similarly? In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14334–14343, 2024.
- Bhaskar Mitra and Nick Craswell. Query auto-completion for rare prefixes. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management (CIKM)*, pp. 1755–1758. Association for Computing Machinery, 2015.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, and Lama Ahmad et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2024.
- Openverse. Openverse: Openly licensed images, audio and more. <https://openverse.org/>, 2025. Accessed: 2025.
- Giacomo Pacini, Fabio Carrara, Nicola Messina, Nicola Tonello, Giuseppe Amato, and Fabrizio Falchi. Maybe you are looking for croqs: Cross-modal query suggestion for text-to-image retrieval. *arXiv preprint arXiv:2412.13834*, 2024.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Christoph Schuhmann. Improved aesthetic predictor: Clip + mlp aesthetic score predictor. <https://github.com/christophschuhmann/improved-aesthetic-predictor>, 2022. GitHub repository.
- Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade W Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa R Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5b: An open large-scale dataset for training next generation image-text models. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022.
- Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15638–15650, 2022.
- Andy Sun, Tianqi Zheng, Aakash Kolekar, Rohit Patki, Hossein Khazaei, Xuan Guo, George Cai, David Liu, Ruirui Li, Yupin Huang, Dante Everaert, Hanqing Lu, Garima Patel, and Monica Cheng. A product-aware query auto-completion framework for e-commerce search via retrieval-augmented generation method. *SIGIR 2024 Workshop on Information Retrieval's Role in RAG Systems (IR-RAG)*, 2024.
- Hailing Wang, Qiaoyu Tian, Liang Li, and Xiaojie Guo. Image demoiréing with a dual-domain distilling network. In *2021 IEEE International Conference on Multimedia and Expo (ICME)*, pp. 1–6, 2021.
- Hailing Wang, Wei Li, Yuanyuan Xi, Jie Hu, Hanting Chen, Longyu Li, and Yunhe Wang. Ift: Image fusion transformer for ghost-free high dynamic range imaging. *arXiv preprint arXiv:2309.15019*, 2023a.

- Hailing Wang, Jianglin Lu, Yitian Zhang, and Yun Fu. Outlier-aware post-training quantization for image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 16175–16184, October 2025.
- Liang Wang, Nan Yang, and Furu Wei. Query2doc: Query expansion with large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 9414–9423. Association for Computational Linguistics, December 2023b.
- Di Wu, Yixin Wan, and Kai-Wei Chang. Visret: Visualization improves knowledge-intensive text-to-image retrieval. *arXiv preprint arXiv:2505.20291*, 2025.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2025.
- Jinyu Yang, Jiali Duan, Son Tran, Yi Xu, Sampath Chanda, Liqun Chen, Belinda Zeng, Trishul Chilimbi, and Junzhou Huang. Vision-language pre-training with triple contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15671–15680, 2022.
- Ran Yi, Haoyuan Tian, Zhihao Gu, Yu-Kun Lai, and Paul L Rosin. Towards artistic image aesthetics assessment: a large-scale dataset and a new method. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 22388–22397, 2023.
- Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022.
- Huimin Zeng, Jiacheng Li, Ziqiang Zheng, and Zhiwei Xiong. All-in-one image compression and restoration. In *Proceedings of the Winter Conference on Applications of Computer Vision (WACV)*, pp. 609–619, February 2025.
- Hongyi Zhu, Jia-Hong Huang, Stevan Rudinac, and Evangelos Kanoulas. Enhancing interactive image retrieval with query rewriting using large language models and vision language models. In *Proceedings of the 2024 International Conference on Multimedia Retrieval*, pp. 978–987, 2024.

## A APPENDIX

### A.1 PROOFS

*Proof of Lemma 1.* Since  $\text{col}(\mathbf{X}_I) = \mathcal{U}$ , we have  $\mathbf{I}_r = \mathbf{X}_I \mathbf{X}_I^\dagger$ . Hence

$$\mathbf{X}_I + \mathbf{P}\mathbf{Y}_I = \mathbf{X}_I + \mathbf{X}_I \mathbf{X}_I^\dagger \mathbf{P}\mathbf{Y}_I = \mathbf{X}_I (\mathbf{I}_r + \mathbf{X}_I^\dagger \mathbf{P}\mathbf{Y}_I). \quad (6)$$

This factorization assumes  $\mathbf{P}\mathbf{Y}_I$  lies in  $\text{col}(\mathbf{X}_I)$ , which is true because  $\mathbf{P}$  projects onto  $\text{col}(\mathbf{X})$  and  $\text{col}(\mathbf{X}) = \text{col}(\mathbf{X}_I)$ . Considering the Neumann Series argument: for any matrix  $\mathbf{E}$ , if  $\|\mathbf{E}\|_2 < 1$ , then  $(\mathbf{I} + \mathbf{E})$  is invertible (i.e., full rank). Given  $\|\mathbf{X}_I^\dagger \mathbf{P}\mathbf{Y}_I\|_2 < 1$ ,  $(\mathbf{I}_r + \mathbf{X}_I^\dagger \mathbf{P}\mathbf{Y}_I)$  is invertible. Thus,  $\text{rank}(\mathbf{X}_I + \mathbf{P}\mathbf{Y}_I) = \text{rank}(\mathbf{X}_I) = r$ .  $\square$

*Proof of Proposition 1.* Since right-multiplication by the orthogonal matrix  $\mathbf{V} = [\mathbf{V}_S \ \mathbf{V}_\perp]$  is rank-preserving, we analyze the following matrices:

$$\mathbf{A}' := \mathbf{A}\mathbf{V} = \mathbf{A}[\mathbf{V}_S \ \mathbf{V}_\perp] = [\mathbf{A}\mathbf{V}_S \ \mathbf{A}\mathbf{V}_\perp] = [\mathbf{A}_S \ \mathbf{0}], \quad (7)$$

$$\mathbf{\Delta}' := \mathbf{\Delta}\mathbf{V} = \mathbf{\Delta}[\mathbf{V}_S \ \mathbf{V}_\perp] = [\mathbf{\Delta}\mathbf{V}_S \ \mathbf{\Delta}\mathbf{V}_\perp] = [\mathbf{\Delta}_S \ \mathbf{\Delta}_\perp], \quad (8)$$

$$\mathbf{B}' := \mathbf{B}\mathbf{V} = (\mathbf{A} + \mathbf{\Delta})\mathbf{V} = \mathbf{A}\mathbf{V} + \mathbf{\Delta}\mathbf{V} = \mathbf{A}' + \mathbf{\Delta}' = [\mathbf{A}_S + \mathbf{\Delta}_S \ \mathbf{\Delta}_\perp], \quad (9)$$

$$\mathbf{C}' := \mathbf{C}\mathbf{V} = \mathbf{C}[\mathbf{V}_S \ \mathbf{V}_\perp] = [\mathbf{C}\mathbf{V}_S \ \mathbf{C}\mathbf{V}_\perp] = [\mathbf{C}_S \ \mathbf{C}_\perp]. \quad (10)$$

Then, for the score matrices, we have:

$$\mathbf{S}_A = \mathbf{A}\mathbf{C}'^\top = (\mathbf{A}\mathbf{V})(\mathbf{C}\mathbf{V})^\top = \mathbf{A}'\mathbf{C}'^\top = [\mathbf{A}_S \ \mathbf{0}] \begin{bmatrix} \mathbf{C}_S^\top \\ \mathbf{C}_\perp^\top \end{bmatrix} = \mathbf{A}_S \mathbf{C}_S^\top, \quad (11)$$

$$\mathbf{S}_B = \mathbf{B}\mathbf{C}'^\top = \mathbf{B}'\mathbf{C}'^\top = [\mathbf{A}_S + \mathbf{\Delta}_S \ \mathbf{\Delta}_\perp] \begin{bmatrix} \mathbf{C}_S^\top \\ \mathbf{C}_\perp^\top \end{bmatrix} = \underbrace{(\mathbf{A}_S + \mathbf{\Delta}_S) \mathbf{C}_S^\top}_X + \underbrace{\mathbf{\Delta}_\perp \mathbf{C}_\perp^\top}_Y. \quad (12)$$

By the SVD construction,  $\mathbf{A}_S$  has full column rank  $r$  and  $\sigma_{\min}(\mathbf{A}_S) = \sigma_r(\mathbf{A}) > 0$ . Since  $\mathbf{\Delta}_S = \mathbf{\Delta}\mathbf{V}_S$  and  $\mathbf{V}_S$  is orthogonal (i.e.,  $\|\mathbf{V}_S\|_2 = 1$ ), it follows that

$$\|\mathbf{\Delta}_S\|_2 = \|\mathbf{\Delta}\mathbf{V}_S\|_2 \leq \|\mathbf{\Delta}\|_2. \quad (13)$$

According to assumption (i), we know  $\|\mathbf{\Delta}_S\|_2 \leq \|\mathbf{\Delta}\|_2 < \sigma_r(\mathbf{A}) = \sigma_{\min}(\mathbf{A}_S)$ . The standard minimum-singular-value perturbation argument (or Weyl's inequality in spectral norm form) yields that  $\mathbf{A}_S + \mathbf{\Delta}_S$  remains full column rank  $r$ . Since left multiplication by a full-column-rank matrix does not change rank, it follows that:

$$\text{rank}(\mathbf{X}) = \text{rank}((\mathbf{A}_S + \mathbf{\Delta}_S) \mathbf{C}_S^\top) = \text{rank}(\mathbf{C}_S^\top) = \text{rank}(\mathbf{C}_S), \quad (14)$$

$$\text{rank}(\mathbf{S}_A) = \text{rank}(\mathbf{A}_S \mathbf{C}_S^\top) = \text{rank}(\mathbf{C}_S) = \text{rank}(\mathbf{X}). \quad (15)$$

According to the assumption (ii), there exists a set  $I$  of size  $r$  such that  $\mathbf{X}_I$  is a basis for  $\text{col}(\mathbf{X})$ . Since  $|\mathbf{X}_I| = r$ , the dimension of the column space must be  $r$ . As a result,

$$\text{rank}(\mathbf{S}_A) = \text{rank}(\mathbf{X}) = r. \quad (16)$$

To prove  $\text{rank}(\mathbf{S}_B) \geq \text{rank}(\mathbf{S}_A)$ , we construct a linear operator (matrix)  $\mathbf{T}$  that decouples the column space of  $\mathbf{X}$  from the rest of the space. Let  $\mathbf{P} := \mathbf{P}_X$  be the orthogonal projector onto  $\mathcal{U} := \text{col}(\mathbf{X})$ ,  $\mathbf{Z} := (\mathbf{I} - \mathbf{P})\mathbf{Y}$  be the component of  $\mathbf{Y}$  orthogonal to  $\mathbf{X}$ , and  $\mathbf{P}_{Z_I}$  be the orthogonal projector onto  $\text{col}(\mathbf{Z}_I)$ . We consider the linear operator:

$$\mathbf{T} = \begin{bmatrix} \mathbf{P} \\ (\mathbf{I} - \mathbf{P}_{Z_I})(\mathbf{I} - \mathbf{P}) \end{bmatrix}. \quad (17)$$

Since  $\mathbf{P}$  projects onto  $\text{col}(\mathbf{X})$ ,  $\mathbf{P}\mathbf{X} = \mathbf{X}$ , and  $(\mathbf{I} - \mathbf{P})\mathbf{X} = \mathbf{0}$ . Given  $\mathbf{Z} := (\mathbf{I} - \mathbf{P})\mathbf{Y}$ , we have

$$\mathbf{T}\mathbf{S}_B = \mathbf{T}(\mathbf{X} + \mathbf{Y}) = \begin{bmatrix} \mathbf{P}\mathbf{X} + \mathbf{P}\mathbf{Y} \\ (\mathbf{I} - \mathbf{P}_{Z_I})(\mathbf{I} - \mathbf{P})(\mathbf{X} + \mathbf{Y}) \end{bmatrix} = \begin{bmatrix} \mathbf{X} + \mathbf{P}\mathbf{Y} \\ (\mathbf{I} - \mathbf{P}_{Z_I})\mathbf{Z} \end{bmatrix}. \quad (18)$$

Restricting to  $I \cup K$ , we analyze the submatrix formed by columns indexed by  $I$  (size  $r$ ) and  $K$  (size determined by  $k$ ), which are disjoint:

$$(\mathbf{TS}_B)_{:,I \cup K} = \begin{bmatrix} \mathbf{X}_I + \mathbf{PY}_I & \mathbf{X}_K + \mathbf{PY}_K \\ (\mathbf{I} - \mathbf{P}_{\mathbf{Z}_I})\mathbf{Z}_I & (\mathbf{I} - \mathbf{P}_{\mathbf{Z}_I})\mathbf{Z}_K \end{bmatrix} = \begin{bmatrix} \mathbf{X}_I + \mathbf{PY}_I & \mathbf{X}_K + \mathbf{PY}_K \\ \mathbf{0} & (\mathbf{I} - \mathbf{P}_{\mathbf{Z}_I})\mathbf{Z}_K \end{bmatrix}. \quad (19)$$

Since  $\mathbf{P}_{\mathbf{Z}_I}$  projects onto the column space of  $\mathbf{Z}_I$ , the residual  $(\mathbf{I} - \mathbf{P}_{\mathbf{Z}_I})\mathbf{Z}_I$  is exactly zero. Because the matrix becomes block upper-triangular, its rank is the sum of the ranks of the diagonal blocks.

By Lemma 1, the top-left block has rank  $r$ . By assumption (iv), the bottom-right block has rank  $k \geq 1$ . Thus block-triangular rank additivity yields

$$\text{rank}((\mathbf{TS}_B)_{:,I \cup K}) = r + k. \quad (20)$$

Since the rank of a submatrix (columns  $I \cup K$ ) is a lower bound for the rank of the full matrix, and left multiplication by  $\mathbf{T}$  cannot increase the rank of the original matrix, we have the inequality:

$$\text{rank}(\mathbf{S}_B) \geq \text{rank}(\mathbf{TS}_B) \geq \text{rank}((\mathbf{TS}_B)_{:,I \cup K}). \quad (21)$$

Given  $k \geq 1$  and  $\text{rank}(\mathbf{S}_A) = r$ , we have:

$$\text{rank}(\mathbf{S}_B) \geq r + k > r = \text{rank}(\mathbf{S}_A). \quad (22)$$

Here, we complete the proof.  $\square$

*Remark.* We decompose  $\Delta$  into two parts: one ( $\Delta_S$ ) that lies in the original row space of  $\mathbf{A}$ , and another ( $\Delta_\perp$ ) that introduces directions outside this space. Assumption (i) ensures the in-span perturbation  $\Delta_S$  is not too large (controlled by  $\sigma_r(\mathbf{A})$ ) so the original  $r$  query directions in  $\mathbf{A}$  are not destroyed by completion. Assumption (ii) asserts that we can select  $r$  columns from  $\mathbf{X}$  that span  $\mathcal{U}$ . This fixes a stable  $r$ -dimensional basis for the existing subspace. Assumption (iii) claims that adding the projected perturbation  $\mathbf{PY}_I$  does not reduce the independence of these  $r$  columns. Thus the original  $r$ -dimensional structure is preserved. Assumption (iv) requires that there exist  $k \geq 1$  columns outside  $I$  whose orthogonal components (after removing projections onto both  $\mathcal{U}$  and  $\text{col}(\mathbf{Z}_I)$ ) are linearly independent. These contribute  $k$  genuinely new directions in  $\mathcal{U}^\perp$ . Together, these assumptions ensure that the rank of  $\mathbf{S}_B$  contains at least the  $r$  preserved dimensions from  $\mathcal{U}$  plus the  $k$  fresh orthogonal ones. Consequently,  $\mathbf{S}_B$  can express more independent scoring patterns and has the ability to potentially make finer-grained distinctions.

## A.2 RELATED WORK

### A.2.1 LARGE LANGUAGE MODELS

Large language models (LLMs) are a class of foundation models designed to process, understand, and generate natural language at scale. With fine-tuning and prompting, these models excel across a variety of tasks, including text generation, summarization, reasoning, translation, and coding (Devlin et al., 2019; Radford et al., 2019; Brown et al., 2020). Notable examples, such as LLaMA3 (Grattafiori et al., 2024), GPT-4o (OpenAI et al., 2024), and Qwen2.5 (Yang et al., 2025), contain billions of parameters and are trained on extensive textual datasets. The large-scale pretraining enables them to capture complex contextual, semantic, and syntactic relationships in natural language. In this work, we utilize pretrained LLMs for query modification to tackle the proposed QCR task. By integrating quality information as conditions, our model autonomously learns to generate quality-aware details for query extension. This provides users with multiple visible query suggestions, allowing them to explore diverse retrieval results.

### A.2.2 VISION-LANGUAGE MODELS

Vision-language models (VLMs) have become the de facto foundation for image-text tasks, demonstrating exceptional potential across a variety of applications (Maniparambil et al., 2024; Dong et al., 2025; Huang et al., 2026). Pioneering work such as CLIP (Radford et al., 2021) and ALIGN (Jia et al., 2021) learn directly from raw texts about images by aligning them in a shared embedding space. CoCa (Yu et al., 2022) combines contrastive loss with captioning loss to train an image-text encoder-decoder model, effectively integrating capabilities from both contrastive and generative approaches. Blip2 (Li et al., 2023) bridges the modality gap with a lightweight Q-Former to improve

pretraining efficiency. In this paper, we adopt joint-embedding VLMs like CLIP as foundation models for text-to-image retrieval. Instead of fine-tuning the VLMs on target datasets, we keep them frozen and focus on refining textual queries to achieve both quality improvement and control over the retrieved images. Improving existing VLMs for retrieval quality control is orthogonal to our approach and represents a promising direction for our future research.

### A.2.3 TEXT-TO-IMAGE RETRIEVAL

Text-to-image retrieval aims to identify the most relevant images from a database given a natural language query (Chen et al., 2021; Lu et al., 2022; Pacini et al., 2024; Wu et al., 2025). It plays a critical role in applications such as visual search, e-commerce, and content-based recommendation. Recent advances in VLMs (Radford et al., 2021; Gao et al., 2022; Yu et al., 2022; Jia et al., 2021) have significantly improved performance on this task by learning powerful cross-modal representations. These models map images and texts into a shared embedding space, typically through contrastive learning on web-scaled image-text pairs. However, existing retrieval systems are primarily optimized to return the top-k images that are semantically aligned with the input query. They overlook other crucial dimensions that strongly affect user satisfaction in practical scenarios, such as aesthetic appeal (Yi et al., 2023), interestingness (Gygli et al., 2013; Abdullahu & Grabner, 2024), or low-level visual quality (Wang et al., 2021; 2023a; Zeng et al., 2025). In this work, we advocate for incorporating quality control into retrieval. By allowing users to explicitly influence the quality attributes of the returned results, we enable a more personalized and controllable search experience, moving beyond simple semantic matching toward a more adaptive, user-centric paradigm.

### A.2.4 QUERY COMPLETION

Query completion (QC) aims to extend user short inputs, referred to as *query prefixes*, by generating longer and more informative *query completions*. It is a widely used technique that helps users better articulate their intent and resolve potential query ambiguity. Traditional QC methods rely on factors such as user profiles, query libraries, and prior search history to extend prefixes into query completions, which limits their applicability to unforeseen prefixes (Bar-Yossef & Kraus, 2011; Mitra & Craswell, 2015; Cai & De Rijke, 2016). Recently, several generative approaches have been proposed for query completion with arbitrary prefixes, primarily for text generation and document retrieval tasks (Lee et al., 2021; Wang et al., 2023b; Lei et al., 2024). QC-based methods for text-to-image retrieval remain scarce, with only a few related works. Zhu et al. (2024) enhance interactive image retrieval through query rewriting based on user relevance feedback, while Sun et al. (2024) leverage LLMs to generate product-aware query completions. Pacini et al. (2024) refine text queries using visual feedback to improve text-to-image retrieval performance. However, these approaches primarily focus on query suggestion rather than achieving control over the quality of retrieved images. In contrast, we tailor query completion to enhance retrieval quality, making the first attempt to adapt it to a given search corpus for quality-controllable retrieval.

## A.3 ANALYSIS OF DATASET DEPENDENCY

Our work focuses on text-to-image retrieval, where the goal is to retrieve relevant images from a fixed dataset based on a textual query. This task is inherently dataset-dependent, as the retrieval process relies entirely on the available images within the dataset. Therefore, the query is crucial in this task: the more specific and detailed the query, the easier the retrieval system can match it to the corresponding image. Conversely, short or vague queries make it significantly more difficult for the system to identify the intended image. That is why our proposed QCQC method aims to enrich the original short queries with more specific, quality-aware details. It is notable that these details are not randomly generated. Instead, they are learned directly from the dataset itself by fine-tuning the LLM to fit the captions. As a result, the completed queries remain dataset-dependent and contextually relevant. The additional details are not unnecessary, as they provide essential guidance to the retrieval system, making it accurately identify images with desired quality.

## A.4 ANALYSIS OF SCORE DIFFERENCES

In Tables 3 and 4, the average quality scores across low, medium, and high conditions may appear close for a set of short queries. This is expected due to dataset limitations. As shown in Figure

Table 7: Retrieval quality with five quality levels on CoCa.

$\mathcal{M}$		Relevance (Red $\rightarrow$ Red)				
		VL	L	M	H	VH
Aesthetics (Green $\leftarrow$ Green)	VL	4.581	4.551	4.559	4.579	4.507
		0.355	0.364	0.372	0.376	0.382
	L	4.870	4.792	4.784	4.748	4.718
		0.357	0.363	0.370	0.376	0.383
	M	4.882	4.954	4.863	4.849	4.820
		0.356	0.366	0.371	0.377	0.381
	H	5.054	5.048	5.005	5.019	4.998
		0.355	0.362	0.371	0.370	0.381
	VH	5.159	5.166	5.161	5.135	5.084
		0.352	0.366	0.369	0.373	0.386

1, the relevance scores across both Flickr2.4M and MS-COCO are not uniformly distributed, and images with extremely low or high scores are rare. For instance, on Flickr2.4M, the relevance scores range from 0.252 to 0.562, and the entire span across the whole dataset is only about 0.3 (where the extreme values correspond to two images from different classes). When retrieving images for a single query, the available results often fall within a narrower score range (much smaller than 0.3) because the dataset lacks images at both ends of the quality spectrum (sparsely distributed). For example, if all images retrieved from the query “a dog” have aesthetic scores between 3.8 and 4.7 (due to dataset limitations), even under the “High” condition, the best available image might score 4.7 which lies in a low range (given the whole range : [2.782, 6.961]). But it is still higher than the score of 3.8 under the “Low” condition. Thus, the method is still effectively ranking and retrieving better images within the constraints of the dataset.

Despite this dataset-level constraint that limits the score differences, our method demonstrates effective ranking ability and a consistent, meaningful trend. As shown from left to right in Tables 1 and 2, both the retrieved image scores and their visual appeal improve progressively as the quality condition increases. This pattern is further supported by quantitative results in Tables 3 and 4, where the average quality scores clearly increase across the low, medium, and high conditions. This behavior cannot be reproduced by baseline methods that lack quality consideration in retrieval.

#### A.5 ADDITIONAL EXPERIMENTAL RESULTS

In Table 7, we present the results on the Flickr2.4M dataset across five quality levels using CoCa as the caption model. Tables 8 and 9 presents the quantitative results of our method using GPT2 (Radford et al., 2019) as the backbone. In addition, we provide more qualitative comparison on the two datasets in Tables 10-12. These results further demonstrate that the proposed QCQC method enables controllable retrieval quality. In rare cases, the completed queries may not align with the semantics of the query prefixes. This occurs when the query completion model generates a sentence referencing different objects. Refer to Table 13 for examples of such cases.

#### A.6 LIMITATION

The relevance and aesthetic quality of the retrieved images depend on the reliability of the VLMs and aesthetic evaluation models. If these models are not sufficiently reliable, retrieval performance can be significantly affected. In addition, the model needs to perceive image quality within the datasets to achieve quality control in retrieval. However, the retrieval datasets may sometimes lack the granularity needed to differentiate between high-quality and low-quality images. In some instances, the retrieval database may not contain high-quality or low-quality images that match specific queries.

#### A.7 THE USE OF LARGE LANGUAGE MODELS (LLMs)

In preparing this paper, large language models were used as writing assistants for grammar checking and minor sentence rephrasing. All technical aspects of the work, including the design, implemen-

Table 8: Retrieval quality of various methods on Flickr2. 4M. CoCa and Blip2 are used to generate textual descriptions; **L** (Low), **M** (Medium), and **H** (High) indicate the quality conditions; and Ctrl specifies whether the method enables controllable retrieval over quality. For both average relevance (Ave Rel) and average aesthetics (Ave Aes), higher values indicate better retrieval quality.

Quality	VLM	Aes Cond Rel Cond	<b>L</b> <b>L</b>	<b>L</b> <b>M</b>	<b>L</b> <b>H</b>	<b>M</b> <b>L</b>	<b>M</b> <b>M</b>	<b>M</b> <b>H</b>	<b>H</b> <b>L</b>	<b>H</b> <b>M</b>	<b>H</b> <b>H</b>	Ctrl ?
Prefix	--	Ave Aes Ave Rel	4.735 0.350	4.735 0.350	4.735 0.350	4.735 0.350	4.735 0.350	4.735 0.350	4.735 0.350	4.735 0.350	4.735 0.350	×
LLaMA3	--	Ave Aes Ave Rel	4.730 0.351	4.822 0.351	4.831 0.351	4.823 0.353	4.837 0.351	4.784 0.350	4.798 0.354	4.722 0.354	4.842 0.352	×
GPT-4o	--	Ave Aes Ave Rel	4.359 0.378	4.651 0.361	4.728 0.357	4.712 0.358	4.668 0.360	4.791 0.356	4.791 0.361	4.816 0.357	5.056 0.361	×
PT	--	Ave Aes Ave Rel	4.681 0.351	4.639 0.344	4.673 0.350	4.688 0.350	4.504 0.346	4.654 0.347	4.610 0.349	4.556 0.352	4.692 0.352	×
FT	CoCa	Ave Aes Ave Rel	4.848 0.366	4.818 0.365	4.864 0.367	4.847 0.365	4.827 0.363	4.876 0.366	4.829 0.367	4.896 0.366	4.853 0.368	×
Ours	CoCa	Ave Aes Ave Rel	4.646 0.354	4.674 0.372	4.632 0.382	4.878 0.355	4.921 0.369	4.894 0.386	5.182 0.357	5.095 0.366	5.124 0.385	✓
FT	Blip2	Ave Aes Ave Rel	4.838 0.369	4.674 0.360	4.744 0.369	4.592 0.365	4.599 0.362	4.772 0.365	4.727 0.373	4.749 0.359	4.818 0.368	×
Ours	Blip2	Ave Aes Ave Rel	4.528 0.355	4.560 0.374	4.470 0.393	4.948 0.354	4.946 0.374	4.885 0.391	5.266 0.354	5.160 0.367	5.236 0.387	✓

Table 9: Retrieval quality of various methods on MS-COCO, where **L** (Low), **M** (Medium), and **H** (High) indicate the quality conditions for retrieval, and Ctrl specifies whether the method enables controllable retrieval over image quality. For both average relevance (Ave Rel) and average aesthetics (Ave Aes), higher values indicate better retrieval quality.

Quality	Aes Cond Rel Cond	<b>L</b> <b>L</b>	<b>L</b> <b>M</b>	<b>L</b> <b>H</b>	<b>M</b> <b>L</b>	<b>M</b> <b>M</b>	<b>M</b> <b>H</b>	<b>H</b> <b>L</b>	<b>H</b> <b>M</b>	<b>H</b> <b>H</b>	Ctrl ?
Prefix	Ave Aes Ave Rel	4.817 0.349	4.817 0.349	4.817 0.349	4.817 0.349	4.817 0.349	4.817 0.349	4.817 0.349	4.817 0.349	4.817 0.349	×
LLaMA3	Ave Aes Ave Rel	4.903 0.348	4.891 0.349	4.855 0.347	4.916 0.348	4.875 0.349	4.880 0.347	4.871 0.348	4.858 0.350	4.911 0.344	×
GPT-4o	Ave Aes Ave Rel	4.673 0.371	4.754 0.357	4.686 0.354	4.782 0.360	4.808 0.358	4.880 0.350	4.838 0.359	5.075 0.352	5.048 0.351	×
PT	Ave Aes Ave Rel	4.742 0.347	4.731 0.345	4.855 0.350	4.821 0.349	4.775 0.344	4.854 0.345	4.830 0.351	4.726 0.347	4.847 0.344	×
FT	Ave Aes Ave Rel	4.785 0.369	4.820 0.369	4.866 0.373	4.813 0.373	4.852 0.369	4.888 0.373	4.833 0.367	4.919 0.376	4.960 0.372	×
FT-CoCa	Ave Aes Ave Rel	4.890 0.347	4.889 0.348	4.793 0.356	4.885 0.346	4.939 0.349	4.903 0.352	4.950 0.347	5.004 0.349	4.898 0.351	×
FT-Blip2	Ave Aes Ave Rel	4.776 0.349	4.883 0.351	4.824 0.352	4.914 0.344	4.968 0.349	4.873 0.350	5.039 0.343	4.967 0.349	5.053 0.349	×
Ours	Ave Aes Ave Rel	4.896 0.354	4.809 0.365	4.719 0.385	4.973 0.356	4.879 0.368	4.916 0.387	5.017 0.353	5.020 0.368	5.109 0.391	✓

tation, and verification of experiments and analyses, were carried out by the authors. In our experiments, we include the LLaMA-3 (Grattafiori et al., 2024) and GPT-4o as query completion models (OpenAI et al., 2024) for empirical comparison.

Table 10: Query completions with their retrieved images and quality scores on MS-COCO

Rel: <b>Low</b> , Aes: <b>Low</b>	Rel: <b>Medium</b> , Aes: <b>Medium</b>	Rel: <b>High</b> , Aes: <b>High</b>
 <p><i>a wine glass next to a plate with some fish and veggies on it</i> Aes 4.790, Rel 0.350</p>	 <p><i>a wine glass next to a plate with some meat and vegetables on it</i> Aes 5.209, Rel 0.376</p>	 <p><i>a wine glass and three clocks all set at different times</i> Aes 5.555, Rel 0.417</p>
 <p><i>a toothbrush on a table with a bunch of scissors</i> Aes 4.149, Rel 0.345</p>	 <p><i>a toothbrush that is on down on the counter</i> Aes 4.837, Rel 0.370</p>	 <p><i>a toothbrush with a smiley face sitting on a sink</i> Aes 5.184, Rel 0.412</p>
 <p><i>a backpack and a line of supplies laying out</i> Aes 3.937, Rel 0.355</p>	 <p><i>a backpack some water rocks and plants</i> Aes 5.130, Rel 0.374</p>	 <p><i>a backpack with rollers is sitting unattended in the middle of this forested dirt road</i> Aes 5.231, Rel 0.470</p>

Table 11: Query completions with their retrieved images and quality scores on Flickr2.4M













Rel: <b>Low</b> , Aes: <b>Low</b>	Rel: <b>Medium</b> , Aes: <b>Medium</b>	Rel: <b>High</b> , Aes: <b>High</b>
 <p><i>a train on a track next to a grassy field</i> Aes 4.585, Rel 0.369</p>	 <p><i>a train station with people waiting to board a bus</i> Aes 4.910, Rel 0.380</p>	 <p><i>a train in the desert</i> Aes 5.488, Rel 0.391</p>
 <p><i>a horse drawn carriage on a dirt road</i> Aes 4.718, Rel 0.357</p>	 <p><i>a horse drawn carriage with people on it</i> Aes 5.023, Rel 0.3773</p>	 <p><i>a horse is grazing in a field under a cloudy sky</i> Aes 5.207, Rel 0.390</p>
 <p><i>a truck is parked in front of a building</i> Aes 4.800, Rel 0.331</p>	 <p><i>a truck is parked under a bridge</i> Aes 5.070, Rel 0.370</p>	 <p><i>a truck is parked in front of the Washington Monument</i> Aes 5.335, Rel 0.389</p>
 <p><i>a boat docked in the water next to other boats</i> Aes 3.814, Rel 0.358</p>	 <p><i>a boat is in the water near a castle</i> Aes 4.839, Rel 0.371</p>	 <p><i>a boat on the water with buildings in the background</i> Aes 5.113, Rel 0.396</p>

Table 12: Query completions with their retrieved images and quality scores on MS-COCO


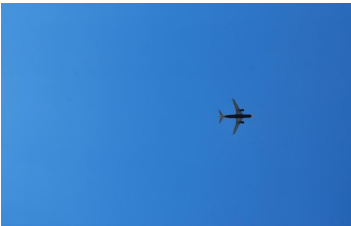










Rel: <b>Low</b> , Aes: <b>Low</b>	Rel: <b>Medium</b> , Aes: <b>Medium</b>	Rel: <b>High</b> , Aes: <b>High</b>
 <p><i>an aeroplane flying in the air with a big blue sky behind it</i> Aes 4.536, Rel 0.354</p>	 <p><i>an aeroplane flying high on a clear sky</i> Aes 4.739, Rel 0.360</p>	 <p><b>Query: an aeroplane flying over the beach and two guys standing on it</b> Aes 5.516, Rel 0.425</p>
 <p><i>a fire hydrant stands in front of a bald eagle wall mural</i> Aes 4.991, Rel 0.355</p>	 <p><i>a fire hydrant sitting in front of a sign for a cafe</i> Aes 5.511, Rel 0.368</p>	 <p><i>a fire hydrant is painted to look like a dalmatian</i> Aes 5.799, Rel 0.447</p>
 <p><i>a toilet with a raised lid in some lavatory</i> Aes 4.457, Rel 0.361</p>	 <p><i>a toilet and sink in a small bathroom with a seat up</i> Aes 4.502, Rel 0.372</p>	 <p><i>a toilet is sitting outside with a sign on it</i> Aes 5.264, Rel 0.394</p>
 <p><i>a sheep is standing on a white fence</i> Aes 4.557, Rel 0.358</p>	 <p><i>a sheep and baby sheep standing in a field</i> Aes 5.014, Rel 0.378</p>	 <p><i>a sheep dog herding sheep through a grass field</i> Aes 5.213, Rel 0.388</p>

Table 13: Some bad retrieval cases on the two datasets

Rel: <b>Low</b> , Aes: <b>Low</b>	Rel: <b>Medium</b> , Aes: <b>Medium</b>	Rel: <b>High</b> , Aes: <b>High</b>
 <p><i>a toilet sits next to a shower an sink</i> Aes 4.192, Rel 0.361</p>	 <p><i>a toilet with a wooden seat on top of it</i> Aes 5.226, Rel 0.371</p>	 <p><i>a toilet in between two trash cans</i> Aes 5.551, Rel 0.434</p>
 <p><i>an apple phone and some other type of machine</i> Aes 3.586, Rel 0.361</p>	 <p><i>an apple and other fruit are sitting together</i> Aes 5.007, Rel 0.373</p>	 <p><i>an apple with a knife stuck into it dripping blood</i> Aes 5.484, Rel 0.395</p>
 <p><i>an orange and blue bath-room with a tub sink and toilet</i> Aes 4.187, Rel 0.357</p>	 <p><i>an orange cat with its eyes closed sitting next to books</i> Aes 4.434, Rel 0.359</p>	 <p><i>an orange and black fire hydrant sitting close to a curb</i> Aes 5.603, Rel 0.393</p>
 <p><i>a clock on the wall of a room</i> Aes 4.578, Rel 0.355</p>	 <p><i>a clock tower with a statue in front of it</i> Aes 5.187, Rel 0.376</p>	 <p><i>a clock tower with a cross on top</i> Aes 5.425, Rel 0.388</p>