

Retrieving Counterfactuals Improves Visual In-Context Learning

Guangzhi Xiong Sanchit Sinha Zhenghao He Aidong Zhang
University of Virginia

Abstract

Vision-language models (VLMs) have achieved impressive performance across a wide range of multimodal reasoning tasks, but they often struggle to disentangle fine-grained visual attributes and reason about underlying causal relationships. In-context learning (ICL) offers a promising avenue for VLMs to adapt to new tasks, but its effectiveness critically depends on the selection of demonstration examples. Existing retrieval-augmented approaches typically rely on passive similarity-based retrieval, which tends to select correlated but non-causal examples, amplifying spurious associations and limiting model robustness. We introduce CIRCLES (Composed Image Retrieval for Causal Learning Example Selection), a novel framework that actively constructs demonstration sets by retrieving counterfactual-style examples through targeted, attribute-guided composed image retrieval. By incorporating counterfactual-style examples, CIRCLES enables VLMs to implicitly reason about the causal relations between attributes and outcomes, moving beyond superficial correlations and fostering more robust and grounded reasoning. Comprehensive experiments on four diverse datasets demonstrate that CIRCLES consistently outperforms existing methods across multiple architectures, especially on small-scale models, with pronounced gains under information scarcity. Furthermore, CIRCLES retrieves more diverse and causally informative examples, providing qualitative insights into how models leverage in-context demonstrations for improved reasoning. Our code is available at <https://github.com/gzxiong/CIRCLES>.

1. Introduction

Vision-language models (VLMs) have recently demonstrated remarkable generalization across a wide spectrum of visual reasoning tasks, achieving impressive results on tasks such as visual question answering [4, 16, 47], image captioning [8, 50], and multimodal classification [1, 14, 38]. This progress is largely attributed to their rich multimodal representations learned through large-scale pre-training. However, in practice, VLMs often struggle with

fine-grained and attribute-sensitive visual reasoning, settings prevalent in real-world vision applications [11, 19, 24].

To address these challenges, in-context learning (ICL) has emerged as a powerful paradigm, enabling VLMs to rapidly adapt to new tasks by conditioning on a small set of demonstration examples provided at test time [3, 43, 51]. Recent work in visual ICL has highlighted the crucial role of retrieval-based selection strategies for constructing effective demonstration sets [9, 12, 28]. Standard approaches, such as RICES (Retrieval-based In-Context Example Selection) [3, 48], assemble in-context examples by identifying the nearest neighbors to the query image within the embedding space. Follow-up works such as MUIER [32] and MMICES [13] augment this process by leveraging multimodal similarity metrics that jointly consider both visual and textual features, enhancing the relevance of retrieved examples to the query.

Despite these advances, such similarity-oriented selection can lead to systematic errors when critical attributes are entangled with irrelevant co-occurrences. Because these retrieval baselines do not explicitly control or intervene on causal factors, models prompted with their selected demonstrations often learn to mimic surface correlations rather than identify the attributes that truly determine the answer. This limitation becomes more pronounced under information scarcity or distribution shifts, where relying solely on correlated neighbors offers little guidance on how changing specific attributes affects the outcome. These observations motivate approaches that go beyond similarity and instead curate demonstration sets that expose disentangled, attribute-level variations.

To address the limitations above, we introduce **CIRCLES** (Composed Image Retrieval for Causal Learning Example Selection), a new framework for visual in-context learning that explicitly incorporates counterfactual reasoning signals into the example selection process. Unlike prior work that selects examples only by image similarity, CIRCLES constructs demonstration sets by explicitly collecting attribute-level counterfactual-style examples via composed image retrieval (CIR), and combining them with standard similarity-based examples to create an informative context.

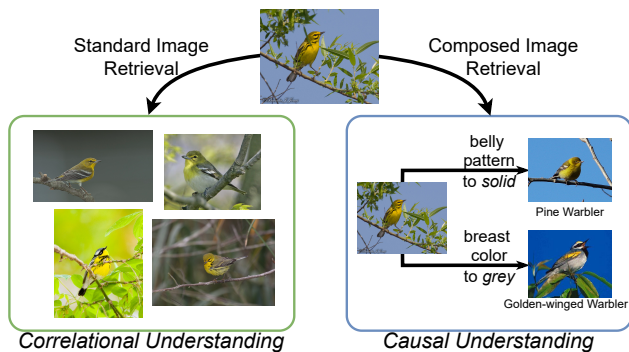


Figure 1. Illustration of how composed image retrieval provides additional causal understanding.

As illustrated in Figure 1, standard similarity-based retrieval provides examples that are visually close to the query but may share irrelevant or confounding attributes. In contrast, by retrieving and composing images that reflect controlled interventions on key attributes (e.g., changing the belly pattern to solid while keeping other attributes as similar as possible), CIR identifies examples that isolate the effect of each attribute on the answer (e.g., similar images with a solid belly pattern are labeled as Pine Warbler), providing more causally informative demonstrations for ICL.

By composing causally informative examples with similarity-based demonstrations, CIRCLES constructs demonstration sets that better reveal the underlying factors driving the correct answer, enabling VLMs to learn disentangled, robust, and interpretable reasoning strategies. Our experiments on four benchmark datasets demonstrate that CIRCLES consistently outperforms existing ICL methods across a range of VLM architectures, with pronounced improvements on small-scale VLMs where models’ internal knowledge is limited. CIRCLES is especially effective under challenging conditions such as information scarcity, where the performance gap widens as relevant data becomes limited. Qualitative analysis further shows that CIRCLES retrieves more diverse and informative examples that clarify critical attributes, providing insights into the ICL process. Our main contributions are as follows:

- We propose CIRCLES, a novel ICL framework that enriches demonstration sets by retrieving counterfactual examples, moving beyond standard similarity-based retrieval.
- Empirical evaluations on multiple image classification and visual question answering datasets demonstrate consistent improvements over existing ICL methods.
- Additional analyses highlight the strength of CIRCLES in low-data regimes and explore optimal practices for constructing informative demonstration sets.

2. Related Work

2.1. Multimodal Reasoning

Vision-language models (VLMs) have rapidly advanced visual tasks such as visual question answering (VQA), image captioning, and multimodal classification by integrating powerful visual encoders with large language models [3, 25, 27]. Despite these gains, several studies have highlighted the limitations of VLMs in compositional and causal reasoning, revealing that state-of-the-art systems often rely on dataset priors or spurious correlations, rather than truly understanding fine-grained attributes or relational structure [2, 22, 39]. For instance, models can achieve high average scores but struggle when asked to reason about the effect of specific attribute changes or to generalize to out-of-distribution queries [21]. These observations motivate methods that structure inference-time evidence around attributes and relations rather than surface similarity.

2.2. Visual In-Context Learning

In-context learning (ICL) enables VLMs to adapt at inference time via conditioning on a small set of demonstration examples [3, 10, 43], but the effectiveness of this process is highly dependent on the structure of the demonstration set [28, 31, 37]. Recent research has shown the sensitivity of VLMs to the order, diversity, and retrieval strategies used to assemble demonstrations [12, 20, 26, 49], demonstrating that naive similarity-based selection can lead to inconsistent performance, sensitivity to confounders, and susceptibility to spurious correlations [7, 18, 52]. These findings motivate moving beyond nearest-neighbor similarity toward example sets that are deliberately diverse and task-aligned. Our approach complements this literature by explicitly introducing counterfactual examples into the demonstration pool, thus enriching the reasoning signal available to the model.

2.3. Composed Image Retrieval

Composed image retrieval (CIR) is a closely related area that focuses on retrieving images matching complex, compositionally specified queries, typically combining a reference image and a manipulation text [15, 40]. Existing CIR techniques are optimized for retrieval accuracy rather than downstream reasoning or ICL composition [5, 44]. Datasets such as FashionIQ [46] and CIRR [29] have advanced research on controlled attribute modifications, but these works treat CIR as an end task. Our use of CIR diverges from this paradigm: instead of retrieving a single modified image, we leverage CIR to construct a set of counterfactually informative demonstrations for in-context learning. By integrating attribute-guided text modifications with VLM-driven captioning, our method operationalizes CIR as a tool for counterfactual intervention, revealing which attributes meaningfully influence downstream predictions.

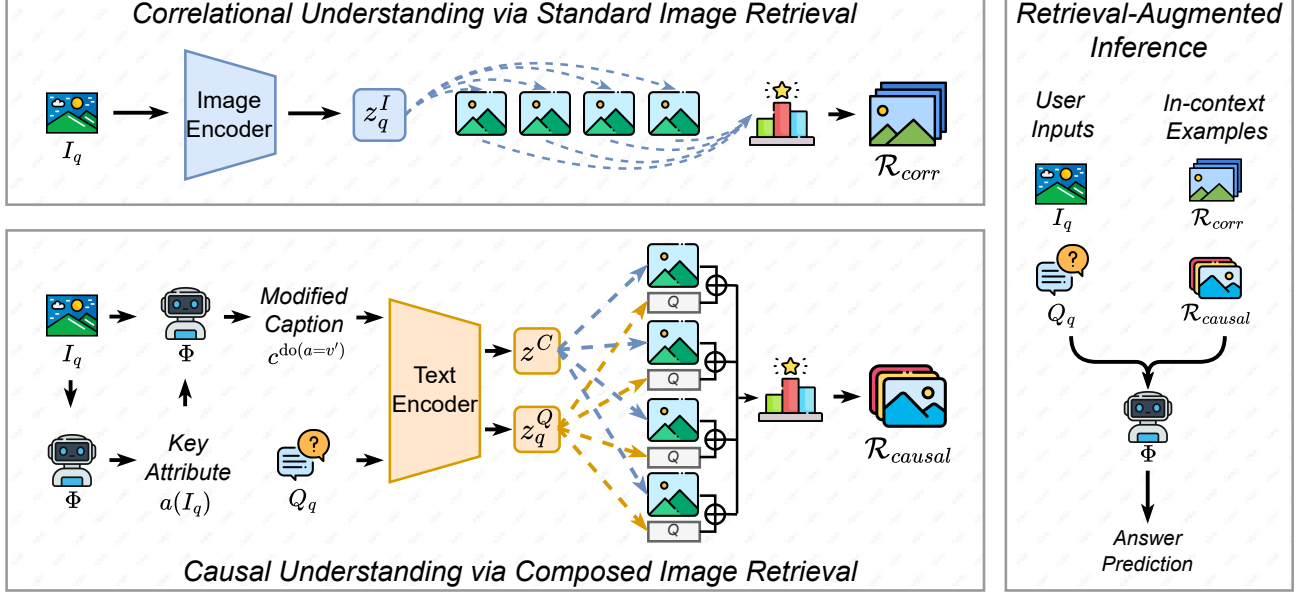


Figure 2. Overview of the CIRCLES framework. Given a query image I_q and question Q_q , the top branch illustrates correlational understanding via standard image retrieval, while the bottom branch depicts causal understanding using attribute-guided composed image retrieval with counterfactual captions from the VLM Φ . Blue/orange rounded rectangles represent image/text embeddings. \mathcal{R}_{corr} and \mathcal{R}_{causal} denote the retrieved in-context examples for answer prediction.

3. Methodology

In this section, we introduce the CIRCLES (Composed Image Retrieval for Causal Learning Example Selection) framework, designed to enhance visual in-context learning for vision-language models (VLMs). CIRCLES does not aim to perform formal causal identification. Instead, it leverages composed image retrieval (CIR) to approximate interventions on key attributes and to facilitate implicit in-context reasoning through contrastive retrieved examples. We begin by formalizing the problem and establishing notation. We then detail the CIRCLES pipeline, which consists of three main components: (1) causal understanding through attribute-guided CIR, (2) correlational understanding via standard image retrieval, and (3) retrieval-augmented inference.

3.1. Task Formulation and Notation

Let $\mathcal{K} = \{(I_j, Q_j, A_j)\}_{j=1}^N$ denote a visual question answering corpus, where I_j is an image, Q_j a natural language question, and A_j the answer. Given a query (I_q, Q_q) , a VLM Φ produces an answer A_q conditioned on the query and an optional retrieved context $\mathcal{R} \subseteq \mathcal{K}$.

We use a frozen CLIP model $E = (f_I, f_T)$ with image encoder f_I and text encoder f_T to obtain L2-normalized embeddings. For each candidate in the training corpus, we precompute embeddings $\mathbf{z}_j^I = f_I(I_j)$ and $\mathbf{z}_j^Q = f_T(Q_j)$. At test time, we compute $\mathbf{z}_q^I = f_I(I_q)$ and $\mathbf{z}_q^Q = f_T(Q_q)$

for the query. An overview of our CIRCLES framework is illustrated in Figure 2.

3.2. Causal Understanding via Attribute-Guided Composed Image Retrieval

To collect demonstrations for causal understanding, CIRCLES identifies semantically meaningful attributes and retrieves counterfactual examples through a two-stage process.

Key Attribute Identification. Given (I_q, Q_q) , we prompt the VLM Φ to extract decisive attribute-value pairs for answering Q_q (e.g., the attribute “breast color” has the value “grey” in Figure 1). Let the set of attributes be

$$\mathcal{A} = \{a_1, \dots, a_m\}, \quad (1)$$

with corresponding values on I_q as

$$\mathbf{v} = (v_1, \dots, v_m), \quad v_i = a_i(I_q). \quad (2)$$

Each attribute a_i has a finite value set \mathcal{V}_i (e.g., possible colors for “breast color”). These attributes form the basis for counterfactual intervention.

Counterfactual Example Retrieval. To isolate the influence of each attribute on the answer, we consider counterfactual variants of the query image by changing one attribute a_i from its original value v_i to an alternative $v'_i \in$

$\mathcal{V}_i \setminus \{v_i\}$ while keeping all other attributes fixed. Direct identification of such counterfactuals in real datasets is generally infeasible. Therefore, for each attribute a_i and candidate counterfactual value v'_i , we prompt the VLM to generate a counterfactual caption $c^{\text{do}(a_i=v'_i)}$ which describes I_q under the atomic intervention $\text{do}(a_i = v'_i)$ [35] with only attribute a_i set to v'_i and all $a_{k \neq i}$ unchanged.

For each candidate example (I_j, Q_j, A_j) in \mathcal{K} , we compute the image-caption similarity score

$$s_j^{\text{img}} = \mathbf{z}_j^{I\top} f_T(c^{\text{do}(a_i=v'_i)}), \quad (3)$$

which measures the visual faithfulness of I_j to the counterfactual scenario described by $c^{\text{do}(a_i=v'_i)}$.

However, relying solely on image-caption similarity can result in retrieving examples that visually align with the counterfactual but are not semantically relevant to the original query, particularly when attribute changes alter the context in unintended ways. To address this, we introduce an additional question-question similarity score

$$s_j^{\text{txt}} = \mathbf{z}_q^{Q\top} \mathbf{z}_j^Q, \quad (4)$$

which serves as a semantic constraint to ensure that retrieved examples not only match the counterfactual attributes but also remain closely related to the original question context.

The score for each candidate is computed as

$$S_j = s_j^{\text{img}} + s_j^{\text{txt}}. \quad (5)$$

By combining both visual faithfulness and semantic relevance, this scoring function identifies training examples that best approximate the intended counterfactual scenario for the given query.

All candidates are ranked according to S_j , and the top- k_{causal} examples are selected to form the set

$$\mathcal{R}_i^{\text{comp}}(v_i \rightarrow v'_i) = \text{TopK}_j(S_j), \quad (6)$$

where TopK_j returns the top examples by the scores. Aggregating across attributes and sampled counterfactuals yields the causal retrieval pool:

$$\mathcal{R}_{\text{causal}} = \bigcup_{i=1}^m \mathcal{R}_i^{\text{comp}}(v_i \rightarrow \tilde{v}'_i), \quad (7)$$

where \tilde{v}'_i is one alternative value sampled from $\mathcal{V}_i \setminus \{v_i\}$ for efficiency, especially when \mathcal{V}_i is large. In practice, we implement this sampling by prompting the VLM to suggest a plausible alternative value for each attribute.

3.3. Correlational Understanding via Standard Image Retrieval

To complement causal understanding given by the counterfactual examples, CIRCLES incorporates a correlation-oriented retrieval stage that provides broader contextual information. Here, we retrieve images that are most similar to

the query image in the embedding space, without enforcing any attribute-based intervention or counterfactual modification. This approach helps the model leverage common visual or semantic patterns present in the dataset, supporting recognition and grounding even when explicit causal factors are absent.

For query image I_q , we compute the image-image similarity score for all candidates in \mathcal{K} as

$$s_j^{\text{corr}} = \mathbf{z}_q^{I\top} \mathbf{z}_j^I, \quad (8)$$

and select the top- k_{corr} most similar examples to form the correlation retrieval set

$$\mathcal{R}_{\text{corr}} = \text{TopK}_j(s_j^{\text{corr}}). \quad (9)$$

While advanced retrieval methods exist [13, 32], we adopt this standard image-only retrieval for its efficiency and generality, which matches the standard RICES-style baseline [3, 48] and cleanly isolates the added value of the proposed causal understanding module. Our experiments in Appendix E further illustrate that image-only retrieval is sufficient when the task is dominated by visual similarity (e.g., classification), while incorporating task text into retrieval becomes important as question semantics become more diverse (e.g., visual question answering).

3.4. Retrieval-Augmented In-Context Learning

The causal and correlational retrieval results are integrated into a unified context for model inference. The final retrieved context is constructed as

$$\mathcal{R} = \mathcal{R}_{\text{causal}} \cup \mathcal{R}_{\text{corr}}, \quad (10)$$

which is used as context for answer generation:

$$A_q = \Phi(I_q, Q_q, \mathcal{R}). \quad (11)$$

By exposing the model to both correlated and counterfactual instances, CIRCLES encourages reasoning that goes beyond surface-level associations and provides deeper insights into the in-context learning process.

4. Experiments

4.1. Experimental Settings

Datasets. For evaluation, we consider a diverse set of datasets: CUB [45] and Flowers [34] for fine-grained image classification, OK-VQA [33] for open-ended visual question answering, and VizWiz [17] for real-world visual question answering under challenging conditions. For CUB and Flowers, we report classification accuracy (Acc) and weighted F1 scores. For OK-VQA and VizWiz, we use exact match (EM) and word-level F1 metrics. These datasets enable a comprehensive assessment of visual reasoning capabilities across various domains.

Table 1. Performance comparison between CIRCLES and baseline methods across different models and datasets. Average scores are computed by treating accuracy (Acc) as EM for classification tasks and averaging EM and F1 across all datasets. The best results for each model and dataset are highlighted in bold.

Model	Method	CUB		Flowers		OK-VQA		VizWiz		Average	
		Acc	F1	Acc	F1	EM	F1	EM	F1	EM	F1
Gemma3 -4B	None	10.56	8.19	46.76	46.50	19.12	26.70	39.85	57.26	29.07	34.66
	Random	9.98	8.10	37.01	39.51	24.81	30.74	51.84	67.08	30.91	36.36
	RICES	65.40	67.62	86.70	87.43	26.65	32.72	56.08	70.40	58.71	64.54
	MUIER	65.21	67.42	86.39	87.28	26.87	32.82	56.59	70.75	58.77	64.57
	MMICES	13.95	10.98	38.09	38.20	26.24	32.41	52.72	68.21	32.75	37.45
	CIRCLES	71.97	72.39	93.32	93.49	31.27	36.89	57.61	71.35	63.54	68.53
Gemma3 -12B	None	30.34	25.02	67.51	64.33	25.47	33.89	56.82	71.26	45.04	48.62
	Random	29.27	25.82	71.28	71.56	33.59	39.56	70.13	80.54	51.07	54.37
	RICES	76.37	76.25	96.44	96.29	36.86	42.61	73.98	83.43	70.91	74.65
	MUIER	76.51	76.39	96.42	96.28	36.58	42.30	73.72	83.04	70.81	74.50
	MMICES	36.37	31.61	71.36	69.13	35.12	40.85	71.50	81.56	53.59	55.79
	CIRCLES	77.03	76.90	97.77	97.75	37.75	43.23	74.30	82.35	71.71	75.06
Qwen2.5 -VL-3B	None	7.09	5.97	22.57	24.04	42.29	46.12	75.83	78.58	36.95	38.68
	Random	15.81	15.63	41.76	44.57	41.54	45.30	74.48	77.40	43.40	45.73
	RICES	72.26	73.88	93.06	92.84	42.57	46.26	70.80	73.37	69.67	71.59
	MUIER	72.25	73.87	93.06	92.83	41.66	45.82	71.61	74.74	69.64	71.81
	MMICES	17.26	15.38	33.84	35.47	39.32	42.95	73.23	76.21	40.91	42.50
	CIRCLES	74.89	76.34	94.70	94.77	43.24	47.27	72.93	75.33	71.44	73.43
Qwen2.5 -VL-7B	None	14.83	11.79	43.47	42.14	33.31	38.33	68.70	73.30	40.08	41.39
	Random	26.27	24.42	47.15	49.74	41.66	46.20	76.11	78.34	47.80	49.67
	RICES	82.15	81.91	98.83	98.85	43.66	48.32	73.79	76.13	74.61	76.30
	MUIER	82.14	81.90	98.93	98.95	44.29	48.76	74.67	77.08	75.01	76.67
	MMICES	34.57	29.72	49.15	47.62	43.40	47.66	70.69	72.77	49.45	49.44
	CIRCLES	82.17	82.13	98.99	99.04	43.54	48.53	77.63	82.36	75.58	78.02

Baselines. We compare CIRCLES against several in-context learning baselines: **None**: zero-shot prompting without in-context examples. **Random**: in-context learning with randomly sampled examples. **RICES** [3, 48]: retrieval based on nearest neighbors using image-image similarity. **MUIER** [32]: multimodal retrieval leveraging both image-image and image-text similarities for example selection. **MMICES** [13]: a two-stage multimodal selector that first retrieves candidates by image-image similarity and then re-ranks them using text-image similarity for improved query alignment. All methods are implemented with Gemma3 (4B/12B) [42] and Qwen2.5-VL (3B/7B) [6] backbones. CLIP (ViT-g/14) [36] is used as the image/text encoder for retrieval, and the in-context example budget is set to 32 for all methods, where CIRCLES has 16 images in $\mathcal{R}_{\text{causal}}$ and 16 in $\mathcal{R}_{\text{corr}}$. More implementation details are provided in Appendix A.

4.2. Performance Comparison

Table 1 presents a comprehensive evaluation of CIRCLES and baseline methods across multiple datasets and vision-language models (VLMs). The results show that all

VLMs benefit from in-context learning, achieving 6.33% to 118.58% relative EM improvements over the zero-shot baseline (None). Additionally, retrieval-based methods (RICES, MUIER, MMICES, CIRCLES) consistently outperform random selection, with up to 105.56% relative improvement, highlighting the importance of selecting relevant examples for effective in-context learning.

Moreover, CIRCLES consistently surpasses other example selection baselines across nearly all datasets and backbone configurations, achieving average relative EM improvements ranging from 0.76% to 94.02%. This highlights the effectiveness of composed image retrieval (CIR) in CIRCLES, which complements the similarity-based retrieval with causal understanding. The performance gains are especially pronounced for smaller backbone models such as Gemma3-4B and Qwen2.5-VL-3B, indicating that CIRCLES provides substantial contextual support when the model’s internal knowledge is limited.

Comparing performance improvements across datasets, we find that CIRCLES yields especially pronounced gains on fine-grained classification tasks (CUB, Flowers), where

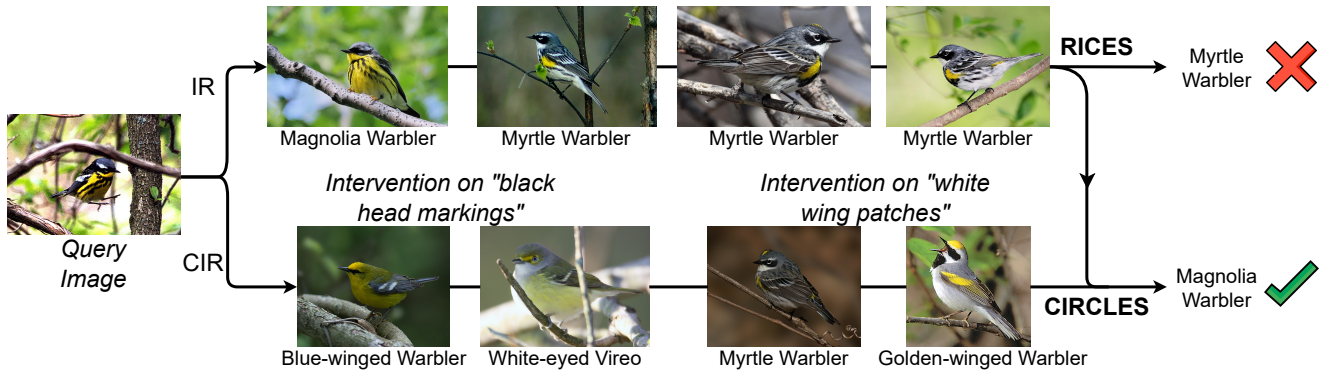


Figure 3. Qualitative comparison of in-context examples retrieved by RICES and CIRCLES for a CUB test image (*Magnolia Warbler*). Top: standard image retrieval (IR) neighbors used by RICES, leading to incorrect predictions. Bottom: counterfactual examples from composed image retrieval (CIR) in CIRCLES, highlighting key attribute changes and guiding the model to the correct label.

distinguishing subtle attribute differences is critical for accurate predictions. This underscores the value of CIR in scenarios requiring nuanced visual reasoning. While improvements on visual question answering tasks (OK-VQA, VizWiz) are more modest, CIRCLES consistently achieves leading results, demonstrating its effectiveness and versatility across diverse vision-language challenges.

4.3. Qualitative Analysis of CIRCLES

Figure 3 provides a qualitative comparison of in-context examples retrieved by RICES (images via standard image retrieval) and CIRCLES (images via both standard and composed image retrieval) for a test sample from the CUB dataset. The query image is a *Magnolia Warbler*, whose label is largely determined by key attributes like “black head markings” and “white wing patches”. RICES, which operates purely in the visual similarity space, retrieves a sequence of images that are globally similar to the query (top row). Because these retrieved examples overrepresent *Myrtle Warblers*, the in-context prompt encourages the model toward the incorrect *Myrtle Warbler* prediction.

In contrast, CIRCLES explicitly constructs counterfactual examples with CIR that intervene on semantically meaningful attributes (bottom row), retrieving images that closely resemble the query but differ in targeted features such as the presence or absence of “black head markings.” This approach retrieves counterfactual examples that clearly demonstrate how changes in key attributes result in different bird species labels. By illustrating class transitions across these counterfactuals, CIRCLES provides the VLM with explicit cues about which attributes are decisive, enabling accurate predictions and offering a more interpretable perspective on the in-context learning process.

4.4. CIRCLES under Information Scarcity

CIRCLES complements standard image retrieval (IR) with CIR to provide diverse and informative in-context exam-

ples, which should be particularly beneficial when the training set contains limited information relevant to the query. To validate this hypothesis, we assess the robustness of CIRCLES and RICES under varying degrees of information scarcity in the training set. Specifically, we simulate different levels of information scarcity by randomly removing a certain percentage of samples from the training set (ranging from 0% to 75%) and evaluate the performance of each method on the CUB dataset using different backbone models. The attribute set used in CIRCLES is fixed across all levels of information scarcity to ensure a fair comparison.

As shown in Figure 4, both RICES and CIRCLES experience performance degradation as more training samples are removed. However, CIRCLES consistently outperforms RICES across all levels of information scarcity. Notably, as the percentage of removed training samples increases, CIRCLES’s performance advantage over RICES becomes more pronounced. In small-scale VLMs such as Gemma3-4B, the relative improvement of CIRCLES over RICES increases from 10.05% to 16.28% when the removal percentage is increased from 0% to 75%. For larger models such as Gemma3-12B, the relative improvements are smaller but still show a consistent upward trend, rising from 0.86% with the full training set to 4.31% when 75% of the samples are removed.

Since CIRCLES shares the same standard image retrieval component as RICES, these results indicate that the CIR component in CIRCLES effectively mitigates the challenges posed by limited relevant information in the training set, enhancing the robustness of in-context learning under information scarcity.

4.5. Impact of CIR Implementation on CIRCLES

To assess the impact of different CIR methods on CIRCLES performance, we compare two representative training-free CIR implementations: CIReVL [23] and OSrCIR [41]. CIReVL generates counterfactual captions by first produc-

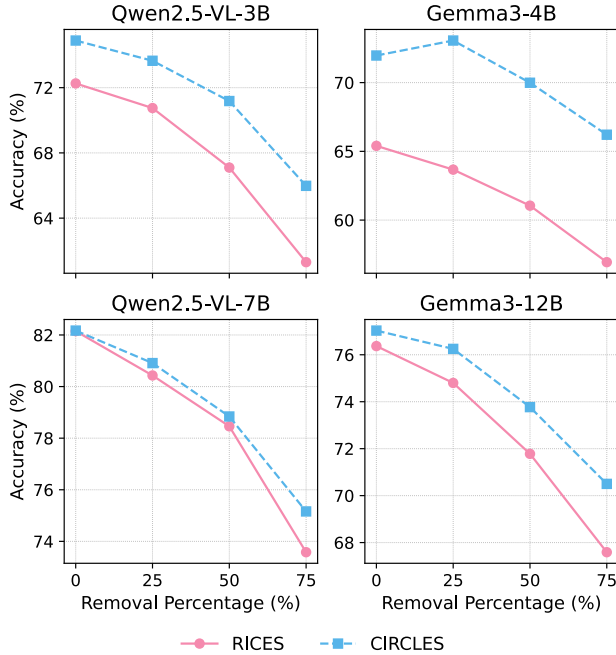


Figure 4. Performance comparison between CIRCLES and RICES on the CUB dataset under varying levels of information scarcity in the training set.

ing a caption for the original image and then editing it to modify target attributes, which can sometimes result in generic or less contextually grounded descriptions. In contrast, OSrCIR directly synthesizes captions conditioned on both the query image and manipulation text, enabling more flexible and fine-grained descriptions of attribute changes and their interactions. This approach allows OSrCIR to capture subtle attribute variations and complex dependencies, producing counterfactuals that are more informative and relevant for in-context learning. In our implementation of CIRCLES, we adopt OSrCIR due to its demonstrated superiority in generating high-quality composed image retrievals [41].

Figure 5 compares CIRCLES instantiated with CIReVL and OSrCIR on the CUB dataset using Gemma3-4B and Gemma3-12B backbones. As shown, CIRCLES with CIR implemented by OSrCIR consistently outperforms CIReVL on both backbone VLMs, achieving relative accuracy improvements ranging from 5.39% to 5.56% and relative F1 improvements from 5.08% to 5.50%. These results align with the CIR performance comparison reported in [41], where OSrCIR demonstrated superior retrieval accuracy over CIReVL. Overall, the quality of the CIR method directly impacts the effectiveness of CIRCLES, as more accurate and contextually relevant composed image retrievals yield better counterfactual examples for in-context learning.

As discussed in Section 3, we further augment CIR with




Figure 5. Comparison of CIRCLES performance using CIR implemented by CIReVL and OSrCIR on the CUB dataset with Gemma3 (4B/12B) backbones.

a question-question similarity component to make the retrieval aware of task similarity (Equation (4)), which is not present in the original CIR formulation [41]. Intuitively, this term biases retrieval toward examples that not only share visual content but also require similar reasoning skills or knowledge types. Because CUB and Flowers have identical questions for all samples, this similarity will not have an effect. Therefore, we evaluate the impact of our proposed enhancement on OK-VQA and VizWiz, where questions are more diverse.

Table 2. Performance comparison of CIRCLES on OK-VQA and VizWiz with (w/) and without (w/o) the question-question similarity component in CIR.

Model	Setting	OK-VQA		VizWiz	
		EM	F1	EM	F1
Gemma-4B	w/o	27.72	33.64	57.40	71.65
	w/	31.27	36.89	57.61	71.35
Gemma-12B	w/o	33.02	39.07	74.37	83.20
	w/	37.75	43.23	74.30	82.35
Qwen-3B	w/o	41.12	45.33	73.28	75.94
	w/	43.24	47.27	72.93	75.33
Qwen-7B	w/o	40.80	46.22	77.29	81.72
	w/	43.54	48.53	77.63	82.36

Table 2 compares CIRCLES performance with and with-

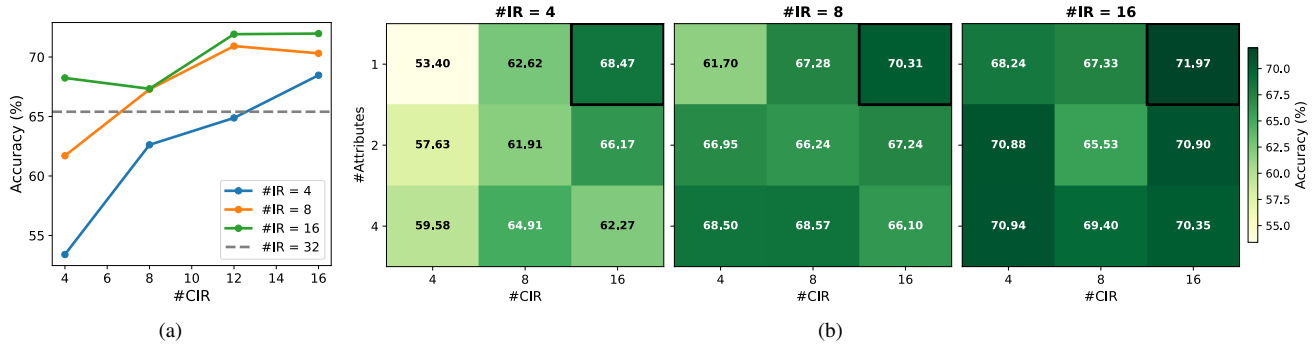


Figure 6. Retrieval budget analysis on CUB with Gemma3-4B. (a) CIRCLES accuracy vs. number of composed images retrieved (#CIR) under different standard retrievals (#IR). (b) Accuracy as a function of the number of intervened attributes (#Attributes) and composed images (#CIR).

out the question-question similarity term. On OK-VQA, we observe consistent improvements across all backbone models, with relative EM gains of +5.16% to 14.32% and F1 gains of +4.28% to 10.65%. This demonstrates that explicitly matching questions enables retrieval of demonstrations with similar reasoning requirements, which is particularly beneficial for knowledge-intensive, open-ended queries. On VizWiz, the improvements are more modest, likely due to the highly diverse and often noisy user-generated questions and challenging images, where textual similarity is a weaker signal and visual cues are more influential. Nevertheless, the addition of the similarity term does not cause significant performance degradation, suggesting it is a robust default that can yield substantial gains when task structures are more regular, as in OK-VQA.

4.6. Retrieval Budget Analysis

We further analyze how to best allocate the in-context example budget between standard image retrieval (IR) and composed image retrieval (CIR) in CIRCLES. Figure 6a shows classification accuracy as a function of the number of composed images retrieved (#CIR) for different settings of standard retrievals (#IR). Increasing #CIR consistently improves accuracy across all #IR configurations, highlighting the benefit of CIR for providing targeted and informative context. Importantly, adding composed images yields substantial gains, even when the total number of retrieved examples is lower than using only standard retrieval (grey line: #CIR=0, #IR=32). This indicates that CIR enables more efficient use of the retrieval budget, as composed examples focus the model on key attribute-level changes and support causal reasoning. Overall, these results demonstrate that combining CIR and IR leads to more informative in-context examples and improved model performance.

While Figure 6a examines the effect of varying the number of composed images (#CIR) when only a single attribute is intervened on (#Attributes = 1), we further analyze how

accuracy changes as both #Attributes and #CIR are varied. Figure 6b presents a heatmap of CIRCLES accuracy on CUB across different hyperparameter settings. When the retrieval budget is limited (e.g., #CIR = 4), distributing interventions across more attributes per query yields better performance than focusing retrievals on a single attribute. As the budget increases (e.g., #CIR = 16), allocating more compositions to fewer attributes becomes advantageous. This suggests that the optimal allocation of the in-context example budget shifts from breadth (intervening on more attributes) to depth (focusing on fewer attributes) as #CIR grows. In practice, these results indicate that under tight budgets, spreading compositions across attributes maximizes coverage, while with larger budgets, concentrating CIR on a smaller subset of attributes provides more precise and informative interventions, because imperfect CIR methods may not always rank the most relevant images at the top. Further discussions on efficiency and additional design ablations are provided in Appendices B, D, and E.

5. Conclusion

We proposed CIRCLES, a retrieval-augmented ICL method that introduces counterfactual demonstrations through attribute-guided composed image retrieval, enabling VLMs to reason beyond surface similarity and better capture the causal structure underlying visual tasks. By integrating compositional interventions with conventional similarity-based retrieval, CIRCLES offers a principled mechanism for enriching demonstration sets with examples that expose disentangled attribute-level variations. Our experiments highlight not only consistent quantitative gains but also qualitative improvements in the diversity and causal informativeness of retrieved examples, particularly under information-scarce regimes where traditional retrieval methods degrade most sharply. These findings suggest that counterfactual retrieval is a practical and effective approach for enhancing visual in-context learning in VLMs.

Acknowledgments

This work is supported in part by the US National Science Foundation (NSF) and the National Institute of Health (NIH) under grants IIS-2106913, IIS-2538206, IIS-2529378, CCF-2217071, CNS-2213700, and R01LM014012-01A1. Any recommendations expressed in this material are those of the authors and do not necessarily reflect the views of NIH or NSF.

References

- [1] Abdelrahman Abdelhamed, Mahmoud Afifi, and Alec Go. What do you see? enhancing zero-shot image classification with multimodal large language models, 2025. [1](#)
- [2] Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Anirudha Kembhavi. Don't just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#)
- [3] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob L Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Saahand Sharifzadeh, Mikołaj Bińkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karén Simonyan. Flamingo: a visual language model for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 23716–23736. Curran Associates, Inc., 2022. [1](#), [2](#), [4](#), [5](#)
- [4] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015. [1](#)
- [5] Muhammad Umer Anwaar, Egor Labintsev, and Martin Kleinsteuber. Compositional learning of image-text query for image retrieval. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1140–1149, 2021. [2](#)
- [6] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhao-hai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report, 2025. [5](#), [1](#)
- [7] Ivana Balazević, David Steiner, Nikhil Parthasarathy, Relja Arandjelović, and Olivier Henaff. Towards in-context scene understanding. In *Advances in Neural Information Processing Systems*, pages 63758–63778. Curran Associates, Inc., 2023. [2](#)
- [8] Davide Bucciarelli, Nicholas Moratelli, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. Personalizing multimodal large language models for image captioning: An experimental analysis. In *Computer Vision – ECCV 2024 Workshops*, pages 351–368, Cham, 2025. Springer Nature Switzerland. [1](#)
- [9] Cheng Chen, Yunpeng Zhai, Yifan Zhao, Jinyang Gao, Bolin Ding, and Jia Li. Provoking multi-modal few-shot lvlm via exploration-exploitation in-context learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3826–3835, 2025. [1](#)
- [10] Huiyi Chen, Jiawei Peng, Kaihua Tang, Xin Geng, and Xu Yang. Enhancing multimodal in-context learning for image classification through coreset optimization. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 5130–5139, 2025. [2](#)
- [11] Meiqi Chen, Bo Peng, Yan Zhang, and Chaochao Lu. CELLO: Causal evaluation of large vision-language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 22353–22374, Miami, Florida, USA, 2024. Association for Computational Linguistics. [1](#)
- [12] Shuo Chen, Zhen Han, Bailan He, Mark Buckley, Philip Torr, Volker Tresp, and Jindong Gu. Understanding and improving in-context learning on vision-language models. In *ICLR 2024 Workshop on Mathematical and Empirical Understanding of Foundation Models*, 2024. [1](#), [2](#)
- [13] Shuo Chen, Zhen Han, Bailan He, Jianzhe Liu, Mark Buckley, Yao Qin, Philip Torr, Volker Tresp, and Jindong Gu. Can multimodal large language models truly perform multimodal in-context learning? In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 6000–6010, 2025. [1](#), [4](#), [5](#)
- [14] Avi Cooper, Keizo Kato, Chia-Hsien Shih, Hiroaki Yamane, Kasper Vinken, Kentaro Takemoto, Taro Sunagawa, Hao-Wei Yeh, Jin Yamanaka, Ian Mason, and Xavier Boix. Rethinking vlms and llms for image classification. *Scientific Reports*, 15(1):19692, 2025. [1](#)
- [15] Longye Du, Shuaiyu Deng, Ying Li, Jun Li, and Qi Tian. A survey on composed image retrieval. *ACM Trans. Multimedia Comput. Commun. Appl.*, 21(7), 2025. [2](#)
- [16] Yifan Du, Junyi Li, Tianyi Tang, Wayne Xin Zhao, and Jirong Wen. Zero-shot visual question answering with language model feedback. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9268–9281, Toronto, Canada, 2023. Association for Computational Linguistics. [1](#)
- [17] Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [4](#)
- [18] Hrayr Harutyunyan, Rafayel Darbinyan, Samvel Karapetyan, and Hrant Khachatryan. In-context learning in presence of spurious correlations. In *ICML 2024 Workshop on In-Context Learning*, 2024. [2](#)
- [19] Cheng-Yu Hsieh, Jieyu Zhang, Zixian Ma, Aniruddha Kembhavi, and Ranjay Krishna. Sugarcrepe: Fixing hackable benchmarks for vision-language compositionality. In *Advances in Neural Information Processing Systems*, pages 31096–31116. Curran Associates, Inc., 2023. [1](#)

- [20] Brandon Huang, Chancharik Mitra, Assaf Arbelle, Leonid Karlinsky, Trevor Darrell, and Roei Herzig. Multimodal task vectors enable many-shot multimodal in-context learning. In *Advances in Neural Information Processing Systems*, pages 22124–22153. Curran Associates, Inc., 2024. 2
- [21] Drew A. Hudson and Christopher D. Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [22] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [23] Shyamgopal Karthik, Karsten Roth, Massimiliano Mancini, and Zeynep Akata. Vision-by-language for training-free compositional image retrieval. In *The Twelfth International Conference on Learning Representations*, 2024. 6
- [24] Jeonghwan Kim and Heng Ji. Finer: Investigating and enhancing fine-grained visual concept recognition in large vision language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6187–6207, Miami, Florida, USA, 2024. Association for Computational Linguistics. 1
- [25] 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 *Proceedings of the 40th International Conference on Machine Learning*, pages 19730–19742. PMLR, 2023. 2
- [26] Li Li, Jiawei Peng, Huiyi Chen, Chongyang Gao, and Xu Yang. How to configure good in-context sequence for visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26710–26720, 2024. 2
- [27] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *Advances in Neural Information Processing Systems*, pages 34892–34916. Curran Associates, Inc., 2023. 2
- [28] Jiachang Liu, Dinghan Shen, Yizhe Zhang, William B Dolan, Lawrence Carin, and Weizhu Chen. What makes good in-context examples for gpt-3? In *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd workshop on knowledge extraction and integration for deep learning architectures*, pages 100–114, 2022. 1, 2
- [29] Zheyuan Liu, Cristian Rodriguez-Opazo, Damien Teney, and Stephen Gould. Image retrieval on real-life images with pre-trained vision-and-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2125–2134, 2021. 2
- [30] Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *Advances in Neural Information Processing Systems*, pages 2507–2521. Curran Associates, Inc., 2022. 5
- [31] Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8086–8098, Dublin, Ireland, 2022. Association for Computational Linguistics. 2
- [32] Yang Luo, Zangwei Zheng, Zirui Zhu, and Yang You. How does the textual information affect the retrieval of multimodal in-context learning? In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5321–5335, Miami, Florida, USA, 2024. Association for Computational Linguistics. 1, 4, 5
- [33] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 4
- [34] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729, 2008. 4
- [35] Judea Pearl. *Causality*. Cambridge University Press, 2 edition, 2009. 4
- [36] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 5, 1
- [37] Ohad Rubin, Jonathan Herzig, and Jonathan Berant. Learning to retrieve prompts for in-context learning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2655–2671, Seattle, United States, 2022. Association for Computational Linguistics. 2
- [38] Oindrila Saha, Grant Van Horn, and Subhransu Maji. Improved zero-shot classification by adapting vlms with text descriptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17542–17552, 2024. 1
- [39] Sanchit Sinha, Guangzhi Xiong, and Aidong Zhang. COCO-tree: Compositional hierarchical concept trees for enhanced reasoning in vision-language models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 2695–2711, Suzhou, China, 2025. Association for Computational Linguistics. 2
- [40] Xuemeng Song, Haoqiang Lin, Haokun Wen, Bohan Hou, Mingzhu Xu, and Liqiang Nie. A comprehensive survey on composed image retrieval. *ACM Trans. Inf. Syst.*, 44(1), 2025. 2
- [41] Yuanmin Tang, Jue Zhang, Xiaoting Qin, Jing Yu, Gaopeng Gou, Gang Xiong, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang, and Qi Wu. Reason-before-retrieve: One-stage reflective chain-of-thoughts for training-free zero-shot composed image retrieval. In *Proceedings of the IEEE/CVF*

- [42] Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, Gaël Liu, Francesco Visin, Kathleen Kenealy, Lucas Beyer, Xiaohai Zhai, Anton Tsitsulin, Robert Busa-Fekete, Alex Feng, Noveen Sachdeva, Benjamin Coleman, Yi Gao, Basil Mustafa, Iain Barr, Emilio Parisotto, David Tian, Matan Eyal, Colin Cherry, Jan-Thorsten Peter, Danila Sinopalnikov, Surya Bhupatiraju, Rishabh Agarwal, Mehran Kazemi, Dan Malkin, Ravin Kumar, David Vilar, Idan Brusilovsky, Jiaming Luo, Andreas Steiner, Abe Friesen, Abhanshu Sharma, Abheesh Sharma, Adi Mayrav Gilady, Adrian Goedeckemeyer, Alaa Saade, Alex Feng, Alexander Kolesnikov, Alexei Bendebury, Alvin Abdagic, Amit Vadi, András György, André Susano Pinto, Anil Das, Ankur Bapna, Antoine Miech, Antoine Yang, Antonia Paterson, Ashish Shenoy, Ayan Chakrabarti, Bilal Piot, Bo Wu, Bobak Shahriari, Bryce Petrin, Charlie Chen, Charline Le Lan, Christopher A. Choquette-Choo, CJ Carey, Cormac Brick, Daniel Deutsch, Danielle Eisenbud, Dee Cattle, Derek Cheng, Dimitris Pappas, Divyashree Shivakumar Sreepathihalli, Doug Reid, Dustin Tran, Dustin Zelle, Eric Noland, Erwin Huizenga, Eugene Kharitonov, Frederick Liu, Gagik Amirkhanyan, Glenn Cameron, Hadi Hashemi, Hanna Klimczak-Plucińska, Harman Singh, Harsh Mehta, Harshal Tushar Lehri, Hussein Hazimeh, Ian Ballantyne, Idan Szpektor, Ivan Nardini, Jean Pouget-Abadie, Jetha Chan, Joe Stanton, John Wieting, Jonathan Lai, Jordi Orbay, Joseph Fernandez, Josh Newlan, Ju yeong Ji, Jyotinder Singh, Kat Black, Kathy Yu, Kevin Hui, Kiran Vodrahalli, Klaus Greff, Linhai Qiu, Marcella Valentine, Marina Coelho, Marvin Ritter, Matt Hoffman, Matthew Watson, Mayank Chaturvedi, Michael Moynihan, Min Ma, Nabila Babar, Natasha Noy, Nathan Byrd, Nick Roy, Nikola Momchev, Nilay Chauhan, Noveen Sachdeva, Oskar Bunyan, Pankil Botarda, Paul Caron, Paul Kishan Rubenstein, Phil Culliton, Philipp Schmid, Pier Giuseppe Sessa, Pingmei Xu, Piotr Stanczyk, Pouya Tafti, Rakesh Shivanna, Renjie Wu, Renke Pan, Reza Rokni, Rob Willoughby, Rohith Vallu, Ryan Mullins, Sammy Jerome, Sara Smoot, Sertan Girgin, Shariq Iqbal, Shashir Reddy, Shruti Sheth, Siim Pöder, Sijal Bhatnagar, Sindhu Raghuram Panyam, Sivan Eiger, Susan Zhang, Tianqi Liu, Trevor Yacovone, Tyler Liechty, Uday Kalra, Utku Evci, Vedant Misra, Vincent Roseberry, Vlad Feinberg, Vlad Kolesnikov, Woohyun Han, Woosuk Kwon, Xi Chen, Yinlam Chow, Yuvein Zhu, Zichuan Wei, Zoltan Egyed, Victor Cotruta, Minh Giang, Phoebe Kirk, Anand Rao, Kat Black, Nabila Babar, Jessica Lo, Erica Moreira, Luiz Gustavo Martins, Omar Sanseviero, Lucas Gonzalez, Zach Gleicher, Tris Warkentin, Vahab Mirrokni, Evan Senter, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, Yossi Matias, D. Sculley, Slav Petrov, Noah Fiedel, Noam Shazeer, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Jean-Baptiste Alayrac, Rohan Anil, Dmitry, Lepikhin, Sebastian Borgeaud, Olivier Bachem, Armand Joulin, Alek Andreev, Cassidy Hardin, Robert Dadashi, and Léonard Hussenot. Gemma 3 technical report, 2025. 5, 1
- [43] Maria Tsimpoukelli, Jacob L Menick, Serkan Cabi, S. M. Ali Eslami, Oriol Vinyals, and Felix Hill. Multimodal few-shot learning with frozen language models. In *Advances in Neural Information Processing Systems*, pages 200–212. Curran Associates, Inc., 2021. 1, 2
- [44] Nam Vo, Lu Jiang, Chen Sun, Kevin Murphy, Li-Jia Li, Li Fei-Fei, and James Hays. Composing text and image for image retrieval - an empirical odyssey. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [45] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 4
- [46] Hui Wu, Yupeng Gao, Xiaoxiao Guo, Ziad Al-Halah, Steven Rennie, Kristen Grauman, and Rogerio Feris. Fashion iq: A new dataset towards retrieving images by natural language feedback. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11307–11317, 2021. 2
- [47] Qi Wu, Damien Teney, Peng Wang, Chunhua Shen, Anthony Dick, and Anton van den Hengel. Visual question answering: A survey of methods and datasets. *Computer Vision and Image Understanding*, 163:21–40, 2017. Language in Vision. 1
- [48] Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Yumao Lu, Zicheng Liu, and Lijuan Wang. An empirical study of gpt-3 for few-shot knowledge-based vqa. In *Proceedings of the AAAI conference on artificial intelligence*, pages 3081–3089, 2022. 1, 4, 5
- [49] Ellen Yi-Ge, Jiechao Gao, Wei Han, and Wei Zhu. DRUM: Learning demonstration retriever for large MULTi-modal models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*, pages 1051–1063, Vienna, Austria, 2025. Association for Computational Linguistics. 2
- [50] Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. Unified vision-language pre-training for image captioning and vqa. In *Proceedings of the AAAI conference on artificial intelligence*, pages 13041–13049, 2020. 1
- [51] Yucheng Zhou, Xiang Li, Qianning Wang, and Jianbing Shen. Visual in-context learning for large vision-language models. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 15890–15902, Bangkok, Thailand, 2024. Association for Computational Linguistics. 1
- [52] Yuhang Zhou, Paiheng Xu, Xiaoyu Liu, Bang An, Wei Ai, and Furong Huang. Explore spurious correlations at the concept level in language models for text classification. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 478–492, Bangkok, Thailand, 2024. Association for Computational Linguistics. 2