



# No “Zero-Shot” Without Exponential Data: Pretraining Concept Frequency Determines Multimodal Model Performance

Vishaal Udandarao<sup>1,2\*</sup> Ameya Prabhu<sup>1,3\*</sup> Adhiraj Ghosh<sup>1</sup> Yash Sharma<sup>1</sup>  
Philip H.S. Torr<sup>3</sup> Adel Bibi<sup>3</sup> Samuel Albanie<sup>2†</sup> Matthias Bethge<sup>1†</sup>

<sup>1</sup>Tübingen AI Center, University of Tübingen    <sup>2</sup>University of Cambridge

<sup>3</sup>University of Oxford

 [github.com/bethgelab/frequency\\_determines\\_performance](https://github.com/bethgelab/frequency_determines_performance)  
 [huggingface.co/datasets/bethgelab/let-it-wag](https://huggingface.co/datasets/bethgelab/let-it-wag)

## Abstract

Web-crawled pretraining datasets underlie the impressive “zero-shot” evaluation performance of multimodal models, such as CLIP for classification/retrieval and Stable-Diffusion for image generation. However, it is unclear how meaningful the notion of “zero-shot” *generalization* is for such multimodal models, as it is not known to what extent their pretraining datasets encompass the downstream concepts targeted for during “zero-shot” evaluation. In this work, we ask: *How is the performance of multimodal models on downstream concepts influenced by the frequency of these concepts in their pretraining datasets?*

We comprehensively investigate this question across 34 models and five standard pretraining datasets (CC-3M, CC-12M, YFCC-15M, LAION-400M, LAION-Aesthetics), generating over 300GB of data artifacts. We consistently find that, far from exhibiting “zero-shot” generalization, multimodal models require exponentially more data to achieve linear improvements in downstream “zero-shot” performance, following a sample inefficient log-linear scaling trend. This trend persists even when controlling for sample-level similarity between pretraining and downstream datasets [79], and testing on purely synthetic data distributions [51]. Furthermore, upon benchmarking models on long-tailed data sampled based on our analysis, we demonstrate that multimodal models across the board perform poorly. We contribute this long-tail test set as the *Let it Wag!* benchmark to further research in this direction. Taken together, our study reveals an exponential need for training data which implies that the key to “zero-shot” generalization capabilities under large-scale training paradigms remains to be found.

## 1 Introduction

Multimodal models like CLIP [91] and Stable Diffusion [96] have revolutionized performance on downstream tasks—CLIP is now the de-facto standard for “zero-shot” image recognition [133, 72, 126, 48, 132] and image-text retrieval [46, 64, 24, 117, 129], while Stable Diffusion is now the de-facto standard for “zero-shot” text-to-image (T2I) generation [93, 17, 96, 41]. In this work, we investigate this empirical success through the lens of zero-shot generalization [69], which refers to the ability of the model to apply its learned knowledge to new unseen concepts. Accordingly, we ask: *Are current multimodal models truly capable of “zero-shot” generalization?*

To address this, we conducted a comparative analysis involving two main factors: (1) the performance of models across various downstream tasks and (2) the frequency of test concepts within their pretraining datasets. We compiled a comprehensive list of 4,029 concepts<sup>1</sup> from 27 downstream tasks spanning classification, retrieval, and image generation, assessing the performance against these concepts. Our analysis spanned

<sup>\*</sup>equal contribution and <sup>†</sup>equal advising, order decided by a coin flip

<sup>1</sup>class categories for classification tasks, objects in the text captions for retrieval tasks, and objects in the text prompts for generation tasks, see Sec. 2 for more details on how we define concepts.

five large-scale pretraining datasets with different scales, data curation methods and sources (CC-3M [107], CC-12M [27], YFCC-15M [113], LAION-Aesthetics [103], LAION-400M [102]), and evaluated the performance of 10 CLIP models and 24 T2I models, spanning different architectures and parameter scales. We consistently find across all our experiments that, across concepts, the frequency of a concept in the pretraining dataset is *a strong predictor* of the model’s performance on test examples containing that concept. Notably, ***model performance scales linearly as the concept frequency in pretraining data grows exponentially*** *i.e.*, we observe a consistent log-linear scaling trend. We find that this log-linear trend is robust to controlling for correlated factors (similar samples in pretraining and test data [79]) and testing across different concept distributions along with samples generated entirely synthetically [51].

Our findings indicate that the impressive empirical performance of multimodal models like CLIP and Stable Diffusion can be largely attributed to the presence of test concepts within their vast pretraining datasets, thus their reported empirical performance does not constitute “zero-shot” generalization. Quite the contrary, these models require exponentially more data on a concept to linearly improve their performance on tasks pertaining to that concept, highlighting extreme sample inefficiency.

In our analysis, we additionally document the distribution of concepts encountered in pretraining data and find that:

- **Concept Distribution:** Across all pretraining datasets, the distribution of concepts is long-tailed (see Fig. 5 in Sec. 5), which indicates that a large fraction of concepts are rare. However, given the extreme sample inefficiency observed, what is rare is not properly learned during multimodal pretraining.
- **Concept Correlation across Pretraining Datasets:** The distribution of concepts across different pretraining datasets are strongly correlated (see Tab. 4 in Sec. 5), which suggests web crawls yield surprisingly similar concept distributions across different pretraining data curation strategies, necessitating explicit rebalancing efforts [11, 125].
- **Image-Text Misalignment between Concepts in Pretraining Data:** Concepts often appear in one modality but not the other, which implies significant misalignment (see Tab. 3 in Sec. 5). Our released data artifacts can help image-text alignment efforts at scale by precisely indicating the examples in which modalities misalign. Note that the log-linear trend across both modalities is robust to this misalignment.

To provide a simple benchmark for generalization performance for multimodal models, which controls for the concept frequency in the training set, we introduce a new long-tailed test dataset called “*Let It Wag!*”. Current models trained on both openly available datasets (*e.g.*, LAION-2B [103], DataComp-1B [46]) and closed-source datasets (*e.g.*, OpenAI-WIT [91], WebLI [29]) have significant drops in performance, providing evidence that our observations may also transfer to closed-source datasets. We publicly release all our data artifacts (over 300GB), amortising the cost of analyzing the pretraining datasets of multimodal foundation models for a more data-centric understanding of the properties of multimodal models in the future.

Several prior works [91, 46, 82, 42, 83, 74] have investigated the role of pretraining data in affecting performance. Mayilvahanan et al. [79] showed that CLIP’s performance is correlated with the similarity between training and test datasets. In other studies on specific areas like question-answering [62] and numerical reasoning [94] in large language models, high train-test set similarity did not fully account for observed performance levels [127]. Our comprehensive analysis of several pretraining image-text datasets significantly adds to this line of work, by (1) showing that concept frequency determines zero-shot performance and (2) pinpointing the exponential need for training data as a fundamental issue for current large-scale multimodal models. We conclude that the key to “zero-shot” generalization capabilities under large-scale training paradigms remains to be found.

## 2 Concepts in Pretraining Data and Quantifying Frequency

In this section, we outline our methodology for obtaining concept frequencies within pretraining datasets. We first define our concepts of interest, then describe algorithms for extracting their frequencies from images

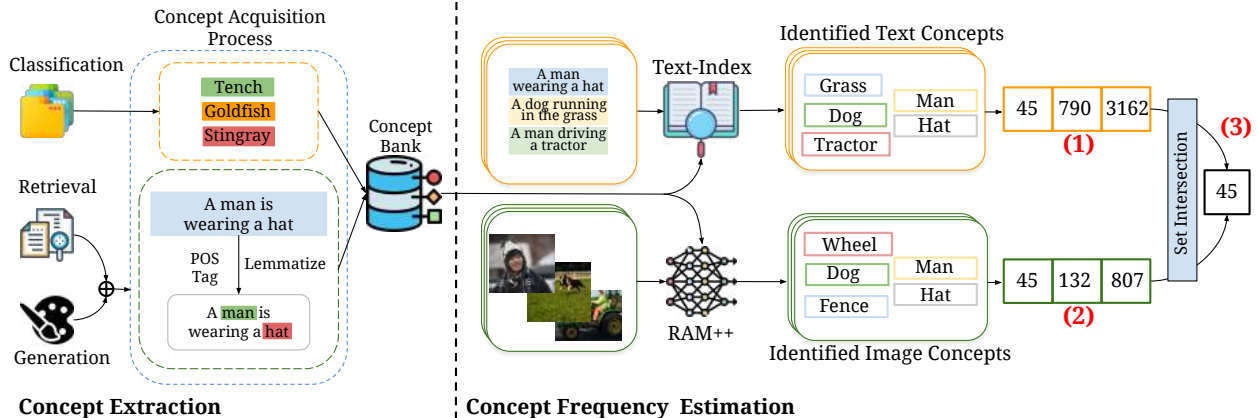


Figure 1: **Concept Extraction and Frequency Estimation Pipeline.** (left) We compile 4,029 concepts from 17 classification, 2 retrieval, and 8 image generation prompt datasets. (right) We construct efficient indices for both text-search (using standard unigram indexing (1)) and image-search (using RAM++ [59] (2)); intersecting hits from both gives us (3) the image-text matched frequencies per concept.

and text captions of pretraining datasets. Finally, we discuss how to aggregate them to calculate matched image-text concept frequencies. For a schematic overview of our methods, see Fig. 1.

**Defining Concepts.** We define “concepts” as the specific objects or class categories we seek to analyze in the pretraining datasets. For zero-shot classification tasks, these concepts are the class names, such as the 1,000 classes in ImageNet [35] (e.g., “tench”, “goldfish”, “stingray”). For image-text retrieval and image generation tasks, concepts are identified as all nouns present in the test set captions or generation prompts, respectively. For example, in the caption, “A man is wearing a hat”, we extract “man” and “hat” as relevant concepts. We additionally filter out nouns that are present in less than five downstream evaluation samples to remove ambiguous or irrelevant concepts. Across all our experiments, we collate a list of 4,029 concepts sourced from 17 classification, 2 retrieval, and 8 image generation downstream datasets (see Tab. 1 for details).

**Concept Frequency from Text Captions.** To enable efficient concept searches, we pre-index all captions from the pretraining datasets, i.e., construct a mapping from concepts to captions. We first use part-of-speech tagging to isolate common and proper nouns and subsequently lemmatize them to standardize word forms [65] with SpaCy [58]. These lemmatized nouns are then cataloged in inverted unigram dictionaries, with each noun being the key and all the indices in the pretraining data samples containing that noun being its values. To determine the frequency of a concept, particularly those composed of multiple words, we examine the concept’s individual unigrams within these dictionaries. For multi-word expressions, by intersecting the lists of sample indices corresponding to each unigram, we identify the samples that contain all parts of the concept. The frequency of the concept in the text captions is the count of these intersecting sample indices. Our frequency estimation algorithm hence allows scalable  $\mathcal{O}(1)$  search with respect to the number of captions for any given concept in the pretraining dataset captions.

**Concept Frequency from Images.** Unlike text captions, we do not have a finite vocabulary for pre-indexing pretraining images, and thus cannot perform  $\mathcal{O}(1)$  concept lookup. Instead, we collect all the 4,029 downstream concepts and verify their presence in images using a pretrained image tagging model. We tested various open-vocabulary object detectors, image-text matching models and multi-tagging models. We found that RAM++ [59]—an open-set tagging model that tags images based on a predefined list of concepts in a multi-label manner—performs the best. This approach generates a list of pretraining images, each tagged with whether the downstream concepts are present or not, from which we can compute concept frequencies. We provide qualitative examples along with design choice ablations in Appx. F.

**Image-Text Matched Concept Frequencies.** Finally, we combine the frequencies obtained from both text and image searches to calculate *matched image-text frequencies*. This involves identifying pretraining

Table 1: Pretraining and downstream datasets used in Image-Text (CLIP) experiments.

Dataset Type	Datasets					
Pretraining	CC-3M [107]		CC-12M [27]	YFCC-15M [113]	LAION-400M [102]	
Classification-Eval	ImageNet [35]	SUN397 [123]	UCF101 [108]	Caltech101 [44]	EuroSAT [55]	CUB [121]
	Caltech256 [49]	Flowers102 [84]	DTD [31]	Birdsnap [15]	Food101 [20]	Stanford-Cars [66]
	FGVCAircraft [77]	Oxford-Pets [87]	Country211 [91]	CIFAR-10 [67]	CIFAR100 [67]	
Retrieval-Eval	Flickr-1K [128]		COCO-5K [73]			

samples where both the image and its associated caption correspond to the concept. By intersecting the lists from our image and text searches, we determine the count of samples that align in both modalities, offering a comprehensive view of concept representation across the dataset. We note that this step is necessary as we observed significant image-text misalignment between concepts in the pretraining datasets (see Tab. 3), hence captions may not reflect what is present in the image and vice-versa. This behaviour has also been alluded to in prior work investigating pretraining data curation strategies [76, 75, 124, 83]. We provide more detailed analysis on image-text misalignment in Sec. 5.

### 3 Comparing Pretraining Frequency & “Zero-Shot” Performance

Having obtained frequency estimates for our downstream concepts, we now establish the relationship between image-text matched pretraining concept frequencies and zero-shot performance across classification, retrieval, and generation tasks. We first detail our experimental approach and then discuss key results.

#### 3.1 Experimental Setup

We analyze two classes of multimodal models: Image-Text and Text-to-Image. For both, we detail the pretraining and testing datasets, along with their associated evaluation parameters.

##### 3.1.1 Image-Text (CLIP) Models

**Datasets.** Our evaluation consists of 4 pretraining datasets, 2 downstream retrieval datasets, and 17 downstream classification datasets, presented in Tab. 1, covering a broad spectrum of objects, scenes, and fine-grained distinctions.

**Models.** We test CLIP [91] models with both ResNet [53] and Vision Transformer [36] architecture, with ViT-B-16 [81] and RN50 [48, 82] trained on CC-3M and CC-12M, ViT-B-16, RN50, and RN101 [61] trained on YFCC-15M, and ViT-B-16, ViT-B-32, and ViT-L-14 trained on LAION400M [102]. We follow `open_clip` [61], `slip` [81] and `cyclip` [48] for all implementation details.

**Prompting.** For zero-shot classification, we experiment with three prompting strategies: {classname} only, “A photo of a {classname}” and prompt-ensembles [91], which averages over 80 different prompt variations of {classname}. For retrieval, we use the image or the caption as input corresponding to I2T (image-to-text) or T2I (text-to-image) retrieval respectively.

**Metrics.** We compute mean zero-shot classification accuracy for classification tasks [91]. For retrieval, we assess performance using traditional metrics for both text-to-image and image-to-text retrieval tasks [91] (Recall@1, Recall@5, Recall@10).

##### 3.1.2 Text-to-Image Models

**Datasets.** Our pretraining dataset is LAION-Aesthetics [103], with downstream evaluations done on subsampled versions of eight datasets as released by HEIM [71]: CUB200 [121], Daily-DALLE [33],

Table 2: Models used in text-to-image (T2I) experiments.

Category	Models			
Models	M-Vader [14]	DeepFloyd-IF-M [9]	DeepFloyd-IF-L [9]	DeepFloyd-IF-XL [9]
	GigaGAN [63]	DALL·E Mini [34]	DALL·E Mega [34]	Promptist+SD-v1.4 [52]
	Dreamlike-Diffusion-v1.0 [2]	Dreamlike Photoreal v2.0 [3]	OpenJourney-v1 [4]	OpenJourney-v2 [5]
	SD-Safe-Max [96]	SD-Safe-Medium [96]	SD-Safe-Strong [96]	SD-Safe-Weak [96]
	SD-v1.4 [96]	SD-v1.5 [96]	SD-v2-Base [96]	SD-v2-1-base [96]
	Vintedois-Diffusion-v0.1 [7]	minDALL·E [97]	Lexica-SD-v1.5 [1]	Redshift-Diffusion [6]

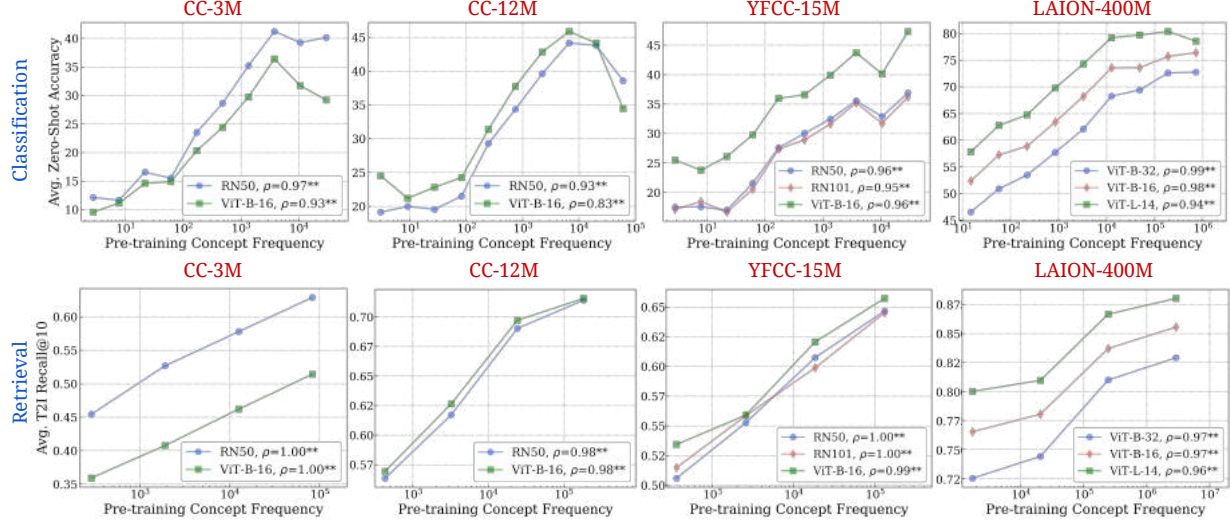


Figure 2: **Log-linear relationships between concept frequency and CLIP zero-shot performance.** Across all tested architectures (RN50, RN101, ViT-B-32, ViT-B-16, ViT-L-14) and pretraining datasets (CC-3M, CC-12M, YFCC-15M, LAION-400M), we observe a consistent linear relationship between CLIP’s zero-shot performance on a concept and the log-scaled concept pretraining frequency. This trend holds for both zero-shot classification (results averaged across 17 datasets) and image-text retrieval (results averaged across 2 datasets). \*\* indicates that the result is significant ( $p < 0.05$  with a two-tailed t-test.), and thus we show pearson correlation ( $\rho$ ) as well.

Detection [30], Parti-Prompts [130], DrawBench [98], COCO-Base [73], Relational Understanding [32] and Winoground [114]. Please refer to HEIM [71] for more details on the evaluation datasets used.

**Models.** We evaluate 24 T2I models, detailed in Tab. 2. Their sizes range from 0.4B parameters (DeepFloyd-IF-M [9] and DALL·E Mini [34]) to 4.3B parameters (DeepFloyd-IF-XL [9]). We include various Stable Diffusion models [96] as well as variants tuned for specific visual styles [6, 4, 5].

**Prompting.** Text prompts from the evaluation datasets are used directly to generate images, with 4 image samples generated for each prompt.

**Metrics.** Evaluation consists of image-text alignment and aesthetic scores. For automated metrics [71], we use expected and max CLIP-score [57] to measure image-text alignment along with expected and max aesthetics-score [102] to measure aesthetics. To verify reliability of the automated metrics, we compare them with human-rated scores (measured on a 5-point grading scale) for both image-text alignment and aesthetics [71]. To supplement the human-rated scores provided by HEIM [71], we confirm our findings by performing a small-scale human evaluation as well (see Appx. C).



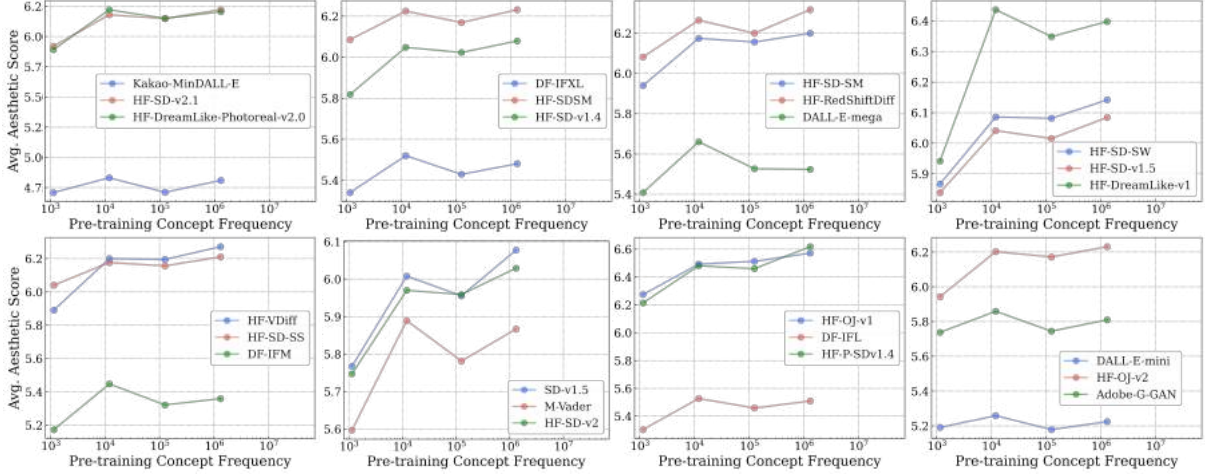


Figure 3: **Log-linear relationships between concept frequency and T2I aesthetic scores.** Across all tested T2I models pretrained on the LAION-Aesthetics dataset, we observe a consistent linear relationship between zero-shot performance on a concept and the log-scaled concept pretraining frequency.

### 3.2 Result: Pretraining Frequency is Predictive of “Zero-Shot” Performance

We now probe the impact of concept frequency in pretraining datasets on the zero-shot performance of image-text models. We utilize the matched image-text concept frequencies for estimating frequency of concepts during pretraining. Our findings, illustrated comprehensively in Figs. 2 and 3, demonstrate the effect concept frequency has on model performance across various tasks and model types.

**Understanding the Plots.** The plots in the main paper present text-image (CLIP) models’ zero-shot classification results using accuracy and text-to-image retrieval performance using Recall@10. Similarly, we present T2I generative models’ performance on image generation tasks using the expected aesthetics score. For the other aforementioned metrics for retrieval as well as other automated generation metrics along with human-rated scores, we find that they show similar trends, and we provide them for reference in Apps. B and C. For clarity, the data presentation is simplified from scatter plots to a cohesive line similar to work from Kandpal et al. [62] and Razeghi et al. [94]. The x-axis is log-scaled, and performance metrics are averaged within bins along this axis for ease-of-visualization of the log-linear correlation. We removed bins containing very few concepts per bin by standard IQR removal [122] following Kandpal et al. [62]. We additionally compute the pearson correlation  $\rho$  for each line and provide significance results based on a two-tailed t-test [110].

**Key Finding: Log-linear scaling between concept frequency and zero-shot performance.** Across all 16 plots, we observe a clear log-linear relationship between concept frequency and zero-shot performance. Note that these plots vary in (i) discriminative vs. generative model types, (ii) classification vs. retrieval tasks, (iii) model architecture and parameter scales, (iv) pretraining datasets with different curation methods and scales, (v) different evaluation metrics, (vi) different prompting strategies for zero-shot classification, and (vii) concept frequencies isolated only from image or text domains (additional experiments which show variation along (v) are presented in Apps. B and C, across (vi) are presented in Appx. A, and across (vii) are presented in Appx. D). The observed log-linear scaling trend persists *across all seven presented dimensions*. Thus, our results clearly reveal data hungry learning, *i.e.*, a lack in current multimodal models’ ability to learn concepts from pretraining datasets in a sample-efficient manner.




Figure 4: **Stress-testing the log-linear scaling trends.** We provide further evidence that the log-linear relationship between performance and concept frequency holds across different scenarios: (*left*) we control for the effect of “similarity” between the downstream test sets and pretraining datasets, and (*right*) we conduct experiments on an entirely synthetic pretraining distribution with no real-world text-captions or images.

## 4 Stress-Testing the Concept Frequency-Performance Scaling Trend

In this section, we seek to isolate the effect of concept frequency on zero-shot performance by controlling a widely known influential factor [127, 79]: similarity in distribution between pretraining and downstream test data. Additionally, we aim to validate our hypothesis further by examining the relationship between concept frequency and downstream performance on models trained on pretraining data with synthetically controlled concept distributions, images and captions.

### 4.1 Controlling for Similar Samples in Pretraining and Downstream Data

**Motivation.** Prior work has suggested that sample-level similarity between pretraining and downstream datasets impacts model performance [62, 79, 127, 94]. This leaves open the possibility that our frequency-performance results are simply an artifact of this factor, *i.e.*, as concept frequency increases, it is likely that the pretraining dataset also contains more similar samples to the test sets. We hence investigate whether concept frequency remains predictive of downstream performance after controlling for sample-level similarity.

**Setup.** We use the LAION-200M [10] dataset for this experiment. We first verified that a CLIP-ViT-B-32 model trained on LAION-200M dataset (used to study sample similarity in prior work [79]) exhibits a similar log-linear trend between concept frequency and zero-shot performance. Then, we use the **near-pruning** method from Mayilvahanan et al. [79] to eliminate 50 million samples most similar to the test sets from the pretraining LAION-200M dataset. We provide details for this in Appx. E.1. This removes the most similar samples between pretraining and test sets. We verify that this procedure influences the performance of the model drastically in performance across our aggregate classification and retrieval tasks respectively, replicating the findings of Mayilvahanan et al. [79].

**Key Finding: Concept Frequency still Predictive of Performance.** We repeat our analysis on models trained with this controlled pretraining dataset with 150M samples, and report results on the same downstream classification and retrieval datasets in Fig. 4 (left). Despite the removal of the most similar samples between pretraining and test sets, we still consistently observe a clear log-linear relationship between pretraining frequency of test set concepts and zero-shot performance.

**Conclusion.** This analysis reaffirms that, despite removing pretraining samples closely related to the test sets, the log-linear relationship between concept frequency and zero-shot performance persists. Note that this is despite substantial decreases in absolute performance, highlighting the robustness of concept frequency as a performance indicator.




Figure 5: **Concept distribution of pre-training datasets is highly long-tailed.** We showcase the distribution of pre-training frequencies of all concepts aggregated across all our downstream classification datasets. Across all three pre-training datasets, we observe very heavy tails. We normalize the concept frequencies and remove concepts with 0 counts for improved readability.

## 4.2 Testing Generalization to Purely Synthetic Concept and Data Distributions

**Motivation.** Sampling across real-world data might not result in significant differences in concept distribution, as we will later show in Sec. 5. Hence, we repeat our analysis on a synthetic dataset designed with an explicitly different concept distribution [51]. This evaluation aims to understand if pretraining concept frequency remains a significant performance predictor within a synthetic concept distribution, generalizing even on models pretrained on entirely synthetic images and captions.

**Setup.** The SynthCI-30M dataset [51] introduces a novel concept distribution, generating 30 million synthetic image-text pairs. Utilizing the publicly available data and models from this benchmark, we explore the relationship between concept frequency and model performance in this synthetic data regime.

**Key Finding: *Concept Frequency is still Predictive of Performance.*** We report results on models trained with their controlled dataset in Fig. 4 (right). We still consistently observe a clear log-linear relationship between concept frequency and zero-shot performance.

**Conclusion.** This consistency highlights that concept frequency is a robust indicator of model performance, extending even to entirely synthetically constructed datasets and pretraining concept distributions.

## 5 Additional Insights from Pretraining Concept Frequencies

We now present notable observations concerning the distribution of downstream concept frequencies across text, image, and text-image matched modalities in pretraining datasets.

**Finding 1: *Pretraining Datasets Exhibit Long-tailed Concept Distribution.*** Our analysis in Fig. 5 reveals an extremely long-tailed distribution of concept frequencies in pretraining datasets, with over two-thirds of concepts occurring at almost negligible frequencies relative to the size of the datasets. Our observations support the findings of past work that have noted the long-tailed distribution of large-scale language datasets [25, 88, 136]. As we observed with the log-linear trend, this distribution directly reflects disparities in performance.

**Finding 2: *Misalignment Between Concepts in Image-Text Pairs.*** We investigated the alignment of concepts within paired pretraining image-text data. Perfect image-text alignment is defined as every image-text pair containing the same concepts. Previous studies have qualitatively discussed the problem of misalignment in large image-text datasets [75, 124, 76]. Our analysis enables us to quantify this *misalignment degree*—for each image-text pair in the pretraining dataset, we find the concepts that are matched to the image and the text caption independently. If there are no intersecting concepts from the independent image



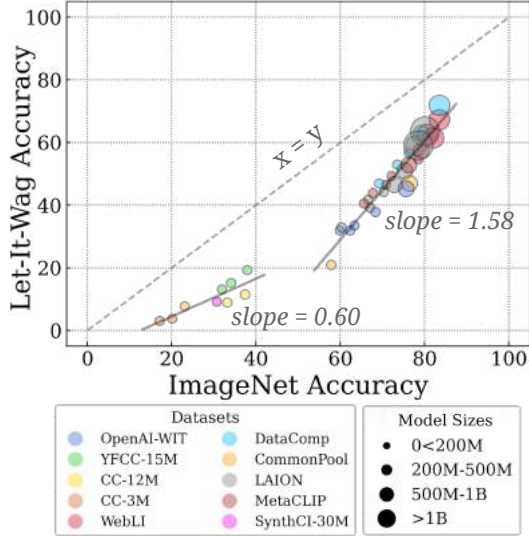


Figure 6: **Large-drops in accuracy on “Let It Wag!”**. Across all 40 tested models, we note large performance drops compared to ImageNet. Further, we note that the gap in performance seems to be decreasing for higher-capacity models as demonstrated by the large positive slope (1.58) for the larger models.

Dataset/ Misalignment	Number of Misaligned pairs	Misalignment Degree (%)
CC3M	557,683	16.81%
CC12M	2,143,784	17.25%
YFCC15M	5,409,248	36.48%
LAION-A	23,104,076	14.34%
LAION400M	21,996,097	5.31%

Table 3: For each pretraining dataset, we present the number of misaligned image-text pairs and the *misalignment degree*: the fraction of misalignment pairs in the dataset.

Correlations	CC3M	CC12M	YFCC15M	L400M
CC3M	1.00	0.79	0.96	0.63
CC12M	–	1.00	0.97	0.74
YFCC15M	–	–	1.00	0.76
L400M	–	–	–	1.00

Table 4: We compute correlation in concept frequency across pretraining datasets. Despite significant differences in scale and curation, we consistently observe strong correlation.

and text hits, we count that pair as misaligned (detailed algorithm provided in Appx. G). Tab. 3 shows the high degree of misalignment in all image-text pairs. To the best of our knowledge, this is the first attempt to explicitly quantify the degree of misalignment in pretraining image-text datasets. We release the precise misaligned image-text pairs in the pretraining datasets to enable better data curation.

**Finding 3: Concept Frequencies Across Datasets are Correlated.** Despite vast differences in the size (ranging from 3M to 400M samples) and curation strategies of the datasets analyzed, we discovered a surprisingly high correlation in concept frequencies across them, as presented in Tab. 4. This consistency suggests that the internet, as the common source of these datasets, naturally exhibits a long-tailed distribution, influencing any dataset derived from it to also display similar long-tailed behavior. This result inspired the “Let It Wag!” dataset.

## 6 Testing the Tail: *Let It Wag!*

**Motivation.** From the previous sections, we have identified a consistent long-tailed concept distribution, highlighting the scarcity of certain concepts on the web. This observation forms the basis of our hypothesis that models are likely to underperform when tested against data distributions that are heavily long-tailed. To test this, we carefully curate 290 concepts that were identified as the least frequent across all pretraining datasets. This includes concepts like an A310 aircraft, a wormsnake, and a tropical kingbird. We then use these concepts to create a classification test set, “Let It Wag!”.

**Dataset Details.** The “Let It Wag!” classification dataset comprises 130K test samples downloaded from the web using the method of Prabhu et al. [90]. The test samples are evenly distributed across 290 categories that represent long-tailed concepts. From the list of curated concepts, we download test set images, deduplicate them, remove outliers, and finally manually clean and hand-verify the labels.

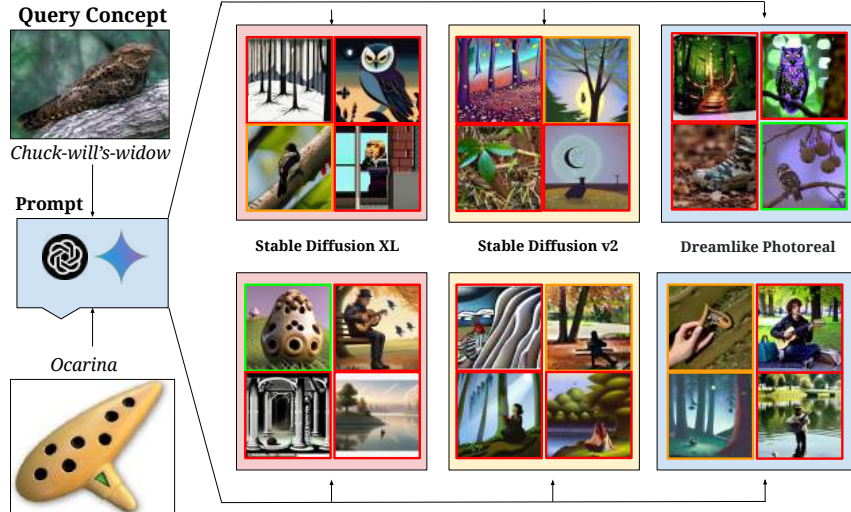


Figure 7: **Qualitative results on the “Let It Wag!” dataset categories demonstrate failure cases of state-of-the-art T2I models on long-tailed concepts.** In our experiments, we create 4 text prompts for each category using Gemini [112] and GPT4 [12] which are fed to 3 Stable Diffusion [96] models. Generation with red border is incorrect, with green border is correct and with yellow border is ambiguous. We observe that despite advances in high-fidelity image generation, there is scope for improvement for such concepts.

**Analysis Details.** We ran both classification and image generation experiments on “Let It Wag!”. For classification, we assessed the performance of 40 text-image (CLIP) models on the “Let It Wag!” classification dataset, using an ensemble of 80 prompts from Radford et al. [91]. For the generative task, we utilized SD-XL [89], SD-v2 [96], and Dreamlike-Photoreal-v2.0 [3] to generate images for the long-tailed concepts. For each model, we ran 50 diffusion steps, maintaining default settings for all other parameters.

**Text-Image Classification Results.** We showcase the results of our long-tailed classification task in Fig. 6—we plot results of all models on both “Let It Wag!” (y-axis) and ImageNet (x-axis). We observe that all models underperform by large margins on the long-tailed “Let It Wag!” dataset (upto 20% lower absolute accuracies compared to ImageNet). This performance drop-off generalises across all model scales and 10 different pretraining data distributions, reinforcing the notion that all web-sourced pretraining datasets are inherently constrained to be long-tailed. With that said, note that the higher capacity models (fitted line with slope=1.58 in Fig. 6) seem to be closing the gap to ImageNet performance, indicating improved performance on the long-tailed concepts.

**T2I Generation Results.** We provide a qualitative analysis on image generation for assessing T2I models on rare concepts in Fig. 7. For diversity, we generate prompts using Gemini [112] (top row of generated images) and GPT4 [12] (bottom row of generated images). Green borders represent correct generations, red borders represent incorrect generations and yellow borders represent ambiguous generation. While descriptive prompting generally aids in improving the quality of generated images [52], we still observe T2I models failing to comprehend and accurately represent many concepts in our “Let It Wag!” dataset. Some failure cases involve misrepresenting activities (such as Pizza Tossing or Cricket Bowling as shown in Fig. 24), generating the wrong concept (Chuck-will’s-widow as shown in Fig. 7 top), as well as not comprehending the concept at all (Ocarina in Fig. 7 bottom). We can see that Stable Diffusion models are prone to the long tail qualitatively—we also provide quantitative results in Appx. H.1.

**Conclusion.** Across both the classification and generation experiments, we have showcased that current multimodal models predictably underperform, regardless of their model scale or pretraining datasets. This suggests a need for better strategies for sample-efficient learning on the long-tail.

## 7 Related Work

**Effect of Pre-training Data on Downstream Data.** Several data-centric prior works [91, 46, 82, 42, 83, 74, 124, 125, 135, 109, 78, 92, 99, 100, 38, 26, 95] have highlighted the importance of pretraining data in affecting performance. Fang et al. [42] robustly demonstrated that pretraining data diversity is the key property underlying CLIP’s strong out-of-distribution generalisation behaviour. Similarly, Berlot-Attwell et al. [16] showed that attribute diversity is crucial for compositional generalization [60], namely systematicity [45]. Nguyen et al. [82] extended the Fang et al. [42] analysis to show that differences in data distributions can predictably change model performance, and that this behaviour can lead to effective data mixing strategies at pretraining time. Mayilvahanan et al. [79] complemented this research direction by showing that CLIP’s performance is correlated with the similarity between training and test datasets. Udandarao et al. [118] further showed that the frequency of certain visual data-types in the LAION-2B dataset was roughly correlated to the performance of CLIP models in identifying visual data-types. Our findings further pinpoint that the frequency of concept occurrences is a key indicator of performance. This complements existing research in specific areas like question-answering [62] and numerical reasoning [94] in large language models, where high train-test set similarity does not fully account for observed performance levels [127]. Concurrent to our work, Parashar et al. [86] also explore the problem of long-tailed concepts in the LAION-2B dataset and how it affects performance of CLIP models, supporting our findings. In contrast to their work, we look at count separately in image and text modalities, as well as across pretraining sets, and do a number of control experiments to thoroughly test the robustness of our result. Finally, our demonstration that the long tail yields a log-linear trend explicitly indicates exponential sample inefficiency in large-scale pretrained models.

**Data-centric analyses.** Our work also adds to the plethora of work that aims to understand and explore the composition of large-scale datasets, and uses data as a medium for improving downstream tasks. Prior work has noted the importance of data for improving model performance on a generalised set of tasks [46, 11, 40, 13, 106]. For instance, several works utilise retrieved and synthetic data for adapting foundation models on a broad set of downstream tasks [119, 54, 115, 21, 101, 134, 90]. Maini et al. [76] observed the existence of “text-centric” clusters in LAION-2B and measured its impact on downstream performance. Other work has sought to target the misalignment problem that we quantified in Tab. 3 by explicit recaptioning of pretraining datasets [68, 28, 120, 131, 83, 17]. Further, studies have also shown that by better data pruning strategies, neural scaling laws can be made more efficient than a power-law [109, 10]. Prior work has also showcased that large-scale datasets suffer from extreme redundancy in concepts, and high degrees of toxic and biased content [39, 116]. Further research has showcased the downstream effects that such biases during pretraining induce in state-of-the-art models [19, 104, 18, 47]. Our work tackles the issue of long-tailed concepts in pretraining datasets, and shows that this is an important research direction to focus efforts on.

## 8 Conclusions and Open Problems

In this work, we delved into the five pretraining datasets of 34 multimodal vision-language models, analyzing the distribution and composition of concepts within, generating over 300GB of data artifacts that we publicly release. Our findings reveal that across concepts, significant improvements in zero-shot performance require exponentially more data, following a log-linear scaling trend. This pattern persists despite controlling for similarities between pretraining and downstream datasets or even when testing models on entirely synthetic data distributions. Further, all tested models consistently underperformed on the “*Let it Wag!*” dataset, which we systematically constructed from our findings to test for long-tail concepts. This underlines a critical reassessment of what “zero-shot” generalization entails for multimodal models, highlighting the limitations in their current generalization capabilities. We highlight a few exciting avenues for future research to bridge these gaps or obtain further insights:

**Understanding Image-Text Misalignments.** One can explore the origins of misalignments between images and texts, such as the limitations of exact matching for concept identification in captions, inaccuracies from the RAM++ tagging model, or captions that are either too noisy or irrelevant.

**Investigating Compositional Generalization.** The term “zero-shot generalization” often refers to models’ ability for compositional generalization, meaning the ability to understand new combinations of concepts not previously encountered. This is distinct from traditional zero-shot learning and presents an intriguing, yet unresolved challenge: analyzing compositional generalization from a data-centric perspective.

**Methods for Bridging the Generalization Gap.** Addressing the challenges posed by the long-tail distribution involves improving model generalization to overcome the limited improvement from pretraining we found in our study. Retrieval mechanisms can compensate for the inherent generalization shortcomings of pretrained models, providing a viable path to mitigating the effects of long-tailed pretraining data distributions.

## Acknowledgements

The authors would like to thank (in alphabetic order): Jonathan Roberts, Karsten Roth, Mehdi Cherti, Prasanna Mayilvahanan, Shyamgopal Karthik and Thao Nguyen for helpful feedback and providing access to various resources throughout the project. YS would like to thank Nicholas Carlini, Daphne Ippolito, Katherine Lee, Matthew Jagielski, and Milad Nasr. AP is funded by Meta AI Grant No. DFR05540. VU and YS thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS). VU also thanks the European Laboratory for Learning and Intelligent Systems (ELLIS) PhD program for support. PT thanks the Royal Academy of Engineering for their support. AB acknowledges the Amazon Research Award. SA is supported by a Newton Trust Grant. MB acknowledges financial support via the Open Philanthropy Foundation funded by the Good Ventures Foundation. This work was supported by the German Research Foundation (DFG): SFB 1233, Robust Vision: Inference Principles and Neural Mechanisms, TP4, project number: 276693517 and the UKRI grant: Turing AI Fellowship EP/W002981/1. MB is a member of the Machine Learning Cluster of Excellence, funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC number 2064/1 – Project number 390727645.

## References

- [1] Lexica search with stable diffusion v1.5 (1b). <https://lexica.art/?q=stable+diffusion+1.5>. 5, 16
- [2] Dreamlike diffusion v1.0. <https://huggingface.co/dreamlike-art/dreamlike-diffusion-1.0>, . 5
- [3] Dreamlike photoreal v2.0. <https://huggingface.co/dreamlike-art/dreamlike-photoreal-2.0>, . 5, 10
- [4] Openjourney v1. <https://huggingface.co/prompthero/openjourney>, . 5
- [5] Openjourney v2. <https://huggingface.co/prompthero/openjourney-v4>, . 5
- [6] Redshift diffusion. <https://huggingface.co/nitrososke/redshift-diffusion>. 5
- [7] Vintedois (22h) diffusion model v0.1. <https://huggingface.co/22h/vintedois-diffusion-v0-1>. 5
- [8] Human (q5). <https://www.wikidata.org/wiki/Q5>. 5
- [9] Deepfloyd if. <https://github.com/deep-floyd/IF>, 2023. 5
- [10] Amro Abbas, Kushal Tirumala, Dániel Simig, Surya Ganguli, and Ari S Morcos. Semdedup: Data-efficient learning at web-scale through semantic deduplication. *Advances in Neural Information Processing Systems*, 2023. 7, 11
- [11] Amro Abbas, Evgenia Rusak, Kushal Tirumala, Wieland Brendel, Kamalika Chaudhuri, and Ari S Morcos. Effective pruning of web-scale datasets based on complexity of concept clusters. In *International Conference on Learning Representations (ICLR)*, 2024. 2, 11
- [12] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 10

- [13] Ekin Akyürek, Tolga Bolukbasi, Frederick Liu, Binbin Xiong, Ian Tenney, Jacob Andreas, and Kelvin Guu. Tracing knowledge in language models back to the training data. In *Findings of the Association for Computational Linguistics: EMNLP*, 2022. 11
- [14] Marco Bellagente, Manuel Brack, Hannah Teufel, Felix Friedrich, Björn Deiseroth, Constantin Eichenberg, Andrew M Dai, Robert Baldock, Souradeep Nanda, Koen Oostermeijer, et al. Multifusion: Fusing pre-trained models for multi-lingual, multi-modal image generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 5
- [15] Thomas Berg, Jiongxin Liu, Seung Woo Lee, Michelle L Alexander, David W Jacobs, and Peter N Belhumeur. Birdsnap: Large-scale fine-grained visual categorization of birds. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 4
- [16] Ian Berlot-Attwell, A Michael Carrell, Kumar Krishna Agrawal, Yash Sharma, and Naomi Saphra. Attribute diversity determines the systematicity gap in vqa. *arXiv preprint arXiv:2311.08695*, 2023. 11
- [17] James Betker, Gabriel Goh, Li Jing, Tim Brooks, Jianfeng Wang, Linjie Li, Long Ouyang, Juntang Zhuang, Joyce Lee, Yufei Guo, et al. Improving image generation with better captions. In *Computer Science*, 2023. 1, 11
- [18] Abeba Birhane and Vinay Uday Prabhu. Large image datasets: A pyrrhic win for computer vision? In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2021. 11
- [19] Abeba Birhane, Sanghyun Han, Vishnu Boddeti, Sasha Luccioni, et al. Into the laion’s den: Investigating hate in multimodal datasets. *Advances in Neural Information Processing Systems*, 2023. 11
- [20] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *European Conference on Computer Vision (ECCV)*, 2014. 4
- [21] Max F Burg, Florian Wenzel, Dominik Zietlow, Max Horn, Osama Makansi, Francesco Locatello, and Chris Russell. Image retrieval outperforms diffusion models on data augmentation. *Transactions on Machine Learning Research (TMLR)*, 2023. 11
- [22] Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramer, and Chiyuan Zhang. Quantifying memorization across neural language models. In *The Eleventh International Conference on Learning Representations*, 2023. 5
- [23] Nicolas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwal, Florian Tramer, Borja Balle, Daphne Ippolito, and Eric Wallace. Extracting training data from diffusion models. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 5253–5270, 2023. 5
- [24] Santiago Castro and Fabian Caba Heilbron. Fitclip: Refining large-scale pretrained image-text models for zero-shot video understanding tasks. *British Machine Vision Conference (BMVC)*, 2022. 1
- [25] Stephanie Chan, Adam Santoro, Andrew Lampinen, Jane Wang, Aaditya Singh, Pierre Richemond, James McClelland, and Felix Hill. Data distributional properties drive emergent in-context learning in transformers. *Conference on Neural Information Processing Systems (NeurIPS)*, 2022. 8
- [26] Kent K Chang, Mackenzie Cramer, Sandeep Soni, and David Bamman. Speak, memory: An archaeology of books known to chatgpt/gpt-4. *arXiv preprint arXiv:2305.00118*, 2023. 11
- [27] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 4, 22
- [28] Lin Chen, Jisong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahua Lin. Sharegpt4v: Improving large multi-modal models with better captions. *arXiv preprint arXiv:2311.12793*, 2023. 11
- [29] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794*, 2022. 2, 22



- [30] Jaemin Cho, Abhay Zala, and Mohit Bansal. Dall-eval: Probing the reasoning skills and social biases of text-to-image generation models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3043–3054, 2023. 5
- [31] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 4
- [32] Colin Conwell and Tomer Ullman. Testing relational understanding in text-guided image generation. *arXiv preprint arXiv:2208.00005*, 2022. 5
- [33] dailydalle2023. Instagram account of daily dall-e. <https://www.instagram.com/dailydall.e/>, 2024. Accessed: 2024-04-03. 4
- [34] Boris Dayma, Suraj Patil, Pedro Cuenca, Khalid Saifullah, Tanishq Abraham, Phuc Le Khac, Luke Melas, and Ritobrata Ghosh. Dall-e mini, 7 2021. URL <https://github.com/borisdyma/dalle-mini>. 5
- [35] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 3, 4
- [36] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 4
- [37] Ling Du, Anthony TS Ho, and Runmin Cong. Perceptual hashing for image authentication: A survey. *Signal Processing: Image Communication*, 2020. 11
- [38] Yanai Elazar, Nora Kassner, Shauli Ravfogel, Amir Feder, Abhilasha Ravichander, Marius Mosbach, Yonatan Belinkov, Hinrich Schütze, and Yoav Goldberg. Measuring causal effects of data statistics on language model’s factual predictions. *arXiv preprint arXiv:2207.14251*, 2022. 11
- [39] Yanai Elazar, Akshita Bhagia, Ian Magnusson, Abhilasha Ravichander, Dustin Schwenk, Alane Suhr, Pete Walsh, Dirk Groeneveld, Luca Soldaini, Sameer Singh, et al. What’s in my big data? *arXiv preprint arXiv:2310.20707*, 2023. 11
- [40] Rahim Entezari, Mitchell Wortsman, Olga Saukh, M Moein Shariatnia, Hanie Sedghi, and Ludwig Schmidt. The role of pre-training data in transfer learning. *arXiv preprint arXiv:2302.13602*, 2023. 11
- [41] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. *arXiv preprint arXiv:2403.03206*, 2024. 1
- [42] Alex Fang, Gabriel Ilharco, Mitchell Wortsman, Yuhao Wan, Vaishaal Shankar, Achal Dave, and Ludwig Schmidt. Data determines distributional robustness in contrastive language image pre-training (clip). *arXiv preprint arXiv:2205.01397*, 2022. 2, 11
- [43] Alex Fang, Albin Madappally Jose, Amit Jain, Ludwig Schmidt, Alexander Toshev, and Vaishaal Shankar. Data filtering networks, 2023. 22
- [44] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *Conference on Computer Vision and Pattern Recognition Workshop (CVPR-W)*, 2004. 4
- [45] Jerry A Fodor and Zenon W Pylyshyn. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1-2):3–71, 1988. 11
- [46] Samir Yitzhak Gadre, Gabriel Ilharco, Alex Fang, Jonathan Hayase, Georgios Smyrnis, Thao Nguyen, Ryan Marten, Mitchell Wortsman, Dhruva Ghosh, Jieyu Zhang, et al. Datacomp: In search of the next generation of multimodal datasets. *arXiv preprint arXiv:2304.14108*, 2023. 1, 2, 11, 22
- [47] Noa Garcia, Yusuke Hirota, Yankun Wu, and Yuta Nakashima. Uncurated image-text datasets: Shedding light on demographic bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6957–6966, 2023. 11

- [48] Shashank Goel, Hritik Bansal, Sumit Bhatia, Ryan Rossi, Vishwa Vinay, and Aditya Grover. Cyclip: Cyclic contrastive language-image pretraining. *arXiv preprint arXiv:2205.14459*, 2022. 1, 4
- [49] Gregory Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. 2007. 4
- [50] Dylan Jasper Hadfield-Menell. *The Principal-Agent Alignment Problem in Artificial Intelligence*. University of California, Berkeley, 2021. 5
- [51] Hasan Abed Al Kader Hammoud, Hani Itani, Fabio Pizzati, Philip Torr, Adel Bibi, and Bernard Ghanem. Synthclip: Are we ready for a fully synthetic clip training? *arXiv preprint arXiv:2402.01832*, 2024. 1, 2, 8, 22
- [52] Yaru Hao, Zewen Chi, Li Dong, and Furu Wei. Optimizing prompts for text-to-image generation. *Advances in Neural Information Processing Systems*, 36, 2024. 5, 10
- [53] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [54] Ruifei He, Shuyang Sun, Xin Yu, Chuhui Xue, Wenqing Zhang, Philip Torr, Song Bai, and Xiaojuan Qi. Is synthetic data from generative models ready for image recognition? *arXiv preprint arXiv:2210.07574*, 2022. 11
- [55] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019. 4
- [56] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. *arXiv preprint arXiv:2104.08718*, 2021. 12
- [57] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning, 2022. 5
- [58] Matthew Honnibal and Ines Montani. spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing. *To appear*, 7(1):411–420, 2017. 3
- [59] Xinyu Huang, Yi-Jie Huang, Youcai Zhang, Weiwei Tian, Rui Feng, Yuejie Zhang, Yanchun Xie, Yaqian Li, and Lei Zhang. Open-set image tagging with multi-grained text supervision. *arXiv e-prints*, pages arXiv–2310, 2023. 3, 12
- [60] Dieuwke Hupkes, Verna Dankers, Mathijs Mul, and Elia Bruni. Compositionality decomposed: How do neural networks generalise? *Journal of Artificial Intelligence Research*, 67:757–795, 2020. 11
- [61] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hannaneh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Openclip, July 2021. URL <https://doi.org/10.5281/zenodo.5143773>. If you use this software, please cite it as below. 4
- [62] Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. Large language models struggle to learn long-tail knowledge. In *International Conference on Machine Learning (ICML)*, pages 15696–15707. PMLR, 2023. 2, 6, 7, 11
- [63] Minguk Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. Scaling up gans for text-to-image synthesis. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 5
- [64] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning (ICML)*, 2021. 1
- [65] Kimmo Koskenniemi. A general computational model for word-form recognition and production. In *10th International Conference on Computational Linguistics and 22nd Annual Meeting of the Association for Computational Linguistics*. The Association for Computational Linguistics, 1984. 3
- [66] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *International Conference on Computer Vision Workshop (ICCV-W)*, 2013. 4

- [67] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. [4](#)
- [68] Zhengfeng Lai, Haotian Zhang, Wentao Wu, Haoping Bai, Aleksei Timofeev, Xianzhi Du, Zhe Gan, Jiulong Shan, Chen-Nee Chuah, Yinfei Yang, et al. From scarcity to efficiency: Improving clip training via visual-enriched captions. *arXiv preprint arXiv:2310.07699*, 2023. [11](#)
- [69] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE transactions on pattern analysis and machine intelligence*, 36(3):453–465, 2013. [1](#)
- [70] Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. Deduplicating training data makes language models better. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8424–8445, 2022. [5](#)
- [71] Tony Lee, Michihiro Yasunaga, Chenlin Meng, Yifan Mai, Joon Sung Park, Agrim Gupta, Yunzhi Zhang, Deepak Narayanan, Hannah Teufel, Marco Bellagente, et al. Holistic evaluation of text-to-image models. *Advances in Neural Information Processing Systems*, 36, 2023. [4](#), [5](#)
- [72] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Conference on Neural Information Processing Systems (NeurIPS)*, 2021. [1](#)
- [73] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision (ECCV)*, 2014. [4](#), [5](#)
- [74] Shayne Longpre, Gregory Yauney, Emily Reif, Katherine Lee, Adam Roberts, Barret Zoph, Denny Zhou, Jason Wei, Kevin Robinson, David Mimno, et al. A pretrainer’s guide to training data: Measuring the effects of data age, domain coverage, quality, & toxicity. *arXiv preprint arXiv:2305.13169*, 2023. [2](#), [11](#)
- [75] Anas Mahmoud, Mostafa Elhoushi, Amro Abbas, Yu Yang, Newsha Ardalani, Hugh Leather, and Ari Morcos. Sieve: Multimodal dataset pruning using image captioning models. *arXiv preprint arXiv:2310.02110*, 2023. [4](#), [8](#)
- [76] Pratyush Maini, Sachin Goyal, Zachary C Lipton, J Zico Kolter, and Aditi Raghunathan. T-mars: Improving visual representations by circumventing text feature learning. *arXiv preprint arXiv:2307.03132*, 2023. [4](#), [8](#), [11](#)
- [77] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. [4](#)
- [78] Daniela Massiceti, Camilla Longden, Agnieszka Slowik, Samuel Wills, Martin Grayson, and Cecily Morrison. Explaining clip’s performance disparities on data from blind/low vision users. *arXiv preprint arXiv:2311.17315*, 2023. [11](#)
- [79] Prasanna Mayilvahanan, Thaddäus Wiedemer, Evgenia Rusak, Matthias Bethge, and Wieland Brendel. Does clip’s generalization performance mainly stem from high train-test similarity? *arXiv preprint arXiv:2310.09562*, 2023. [1](#), [2](#), [7](#), [11](#)
- [80] Matthias Minderer, Alexey Gritsenko, and Neil Houlsby. Scaling open-vocabulary object detection. *Advances in Neural Information Processing Systems*, 36, 2023. [12](#)
- [81] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. *arXiv preprint arXiv:2112.12750*, 2021. [4](#)
- [82] Thao Nguyen, Gabriel Ilharco, Mitchell Wortsman, Sewoong Oh, and Ludwig Schmidt. Quality not quantity: On the interaction between dataset design and robustness of clip. *arXiv preprint arXiv:2208.05516*, 2022. [2](#), [4](#), [11](#)
- [83] Thao Nguyen, Samir Yitzhak Gadre, Gabriel Ilharco, Sewoong Oh, and Ludwig Schmidt. Improving multimodal datasets with image captioning. *Advances in Neural Information Processing Systems*, 2023. [2](#), [4](#), [11](#)
- [84] 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. IEEE, 2008. [4](#)

- [85] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 16
- [86] Shubham Parashar, Zhiqiu Lin, Tian Liu, Xiangjue Dong, Yanan Li, Deva Ramanan, James Caverlee, and Shu Kong. The neglected tails of vision-language models. *CVPR*, 2024. 11
- [87] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 4
- [88] Steven T Piantadosi. Zipf’s word frequency law in natural language: A critical review and future directions. *Psychonomic bulletin & review*, 21:1112–1130, 2014. 8
- [89] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 10, 16
- [90] Ameya Prabhu, Hasan Abed Al Kader Hammoud, Ser-Nam Lim, Bernard Ghanem, Philip HS Torr, and Adel Bibi. From categories to classifier: Name-only continual learning by exploring the web. *arXiv preprint arXiv:2311.11293*, 2023. 9, 11
- [91] 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 (ICML)*, 2021. 1, 2, 4, 10, 11, 22
- [92] Vivek Ramanujan, Thao Nguyen, Sewoong Oh, Ali Farhadi, and Ludwig Schmidt. On the connection between pre-training data diversity and fine-tuning robustness. *Advances in Neural Information Processing Systems*, 36, 2024. 11
- [93] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning (ICML)*, 2021. 1
- [94] Yasaman Razeghi, Robert L Logan IV, Matt Gardner, and Sameer Singh. Impact of pretraining term frequencies on few-shot numerical reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 840–854, 2022. 2, 6, 7, 11
- [95] Yasaman Razeghi, Raja Sekhar Reddy Mekala, Robert L Logan Iv, Matt Gardner, and Sameer Singh. Snoopy: An online interface for exploring the effect of pretraining term frequencies on few-shot lm performance. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 389–395, 2022. 11
- [96] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 5, 10, 16
- [97] Kim Saehoon, Cho Sanghun, Kim Chiheon, Doyup Lee, and Woonhyuk Baek. mindall-e on conceptual captions. <https://github.com/kakaobrain/minDALL-E>, 2021. 5
- [98] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. 5
- [99] Dvir Samuel, Rami Ben-Ari, Simon Raviv, Nir Darshan, and Gal Chechik. Generating images of rare concepts using pretrained diffusion models. *arXiv preprint arXiv:2304.14530*, 18, 2023. 11
- [100] Shibani Santurkar, Yann Dubois, Rohan Taori, Percy Liang, and Tatsunori Hashimoto. Is a caption worth a thousand images? a controlled study for representation learning. *arXiv preprint arXiv:2207.07635*, 2022. 11

- [101] Mert Bülent Sarıyıldız, Kartteek Alahari, Diane Larlus, and Yannis Kalantidis. Fake it till you make it: Learning transferable representations from synthetic imagenet clones. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 11
- [102] 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. 2, 4, 5, 22
- [103] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 2, 4, 22
- [104] Preethi Seshadri, Sameer Singh, and Yanai Elazar. The bias amplification paradox in text-to-image generation. *arXiv preprint arXiv:2308.00755*, 2023. 11
- [105] Hassan Shahmohammadi, Adhiraj Ghosh, and Hendrik Lensch. Vipe: Visualise pretty-much everything. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5477–5494, 2023. 16
- [106] Hanyin Shao, Jie Huang, Shen Zheng, and Kevin Chen-Chuan Chang. Quantifying association capabilities of large language models and its implications on privacy leakage. *arXiv preprint arXiv:2305.12707*, 2023. 11
- [107] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018. 2, 4, 22
- [108] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 4
- [109] Ben Sorscher, Robert Geirhos, Shashank Shekhar, Surya Ganguli, and Ari Morcos. Beyond neural scaling laws: beating power law scaling via data pruning. *Advances in Neural Information Processing Systems*, 35: 19523–19536, 2022. 11
- [110] Student. Probable error of a correlation coefficient. *Biometrika*, pages 302–310, 1908. 6
- [111] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 11
- [112] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. 10
- [113] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016. 2, 4, 22
- [114] Tristan Thrush, Ryan Jiang, Max Bartolo, Amanpreet Singh, Adina Williams, Douwe Kiela, and Candace Ross. Winoground: Probing vision and language models for visio-linguistic compositionality. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 5
- [115] Yonglong Tian, Lijie Fan, Phillip Isola, Huiwen Chang, and Dilip Krishnan. Stablerep: Synthetic images from text-to-image models make strong visual representation learners. *Conference on Neural Information Processing Systems (NeurIPS)*, 2023. 11
- [116] Kushal Tirumala, Daniel Simig, Armen Aghajanyan, and Ari Morcos. D4: Improving llm pretraining via document de-duplication and diversification. *Advances in Neural Information Processing Systems*, 36, 2024. 11
- [117] Vishaal Udandaraao, Abhishek Maiti, Deepak Srivatsav, Suryatej Reddy Vyalla, Yifang Yin, and Rajiv Ratn Shah. Cobra: Contrastive bi-modal representation algorithm. *arXiv preprint arXiv:2005.03687*, 2020. 1



- [118] Vishaal Udandaraao, Max F Burg, Samuel Albanie, and Matthias Bethge. Visual data-type understanding does not emerge from scaling vision-language models. *arXiv preprint arXiv:2310.08577*, 2023. 11
- [119] Vishaal Udandaraao, Ankush Gupta, and Samuel Albanie. Sus-x: Training-free name-only transfer of vision-language models. In *International Conference on Computer Vision (ICCV)*, 2023. 11
- [120] Pavan Kumar Anasosalu Vasu, Hadi Pouransari, Fartash Faghri, Raviteja Vemulapalli, and Oncel Tuzel. Mobileclip: Fast image-text models through multi-modal reinforced training. *arXiv preprint arXiv:2311.17049*, 2023. 11
- [121] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 4
- [122] Steven Walfish. A review of statistical outlier methods. *Pharmaceutical technology*, 30(11):82, 2006. 6
- [123] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010. 4
- [124] Hu Xu, Saining Xie, Po-Yao Huang, Licheng Yu, Russell Howes, Gargi Ghosh, Luke Zettlemoyer, and Christoph Feichtenhofer. Cit: Curation in training for effective vision-language data, 2023. 4, 8, 11
- [125] Hu Xu, Saining Xie, Xiaoqing Ellen Tan, Po-Yao Huang, Russell Howes, Vasu Sharma, Shang-Wen Li, Gargi Ghosh, Luke Zettlemoyer, and Christoph Feichtenhofer. Demystifying clip data, 2023. 2, 11, 22
- [126] 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 *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1
- [127] Gregory Yauney, Emily Reif, and David Mimno. Data similarity is not enough to explain language model performance. *arXiv preprint arXiv:2311.09006*, 2023. 2, 7, 11
- [128] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78, 2014. 4
- [129] 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. 1
- [130] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2(3):5, 2022. 5
- [131] Qiying Yu, Quan Sun, Xiaosong Zhang, Yufeng Cui, Fan Zhang, Xinlong Wang, and Jingjing Liu. Capsfusion: Rethinking image-text data at scale. *arXiv preprint arXiv:2310.20550*, 2023. 11
- [132] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1
- [133] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. *arXiv preprint arXiv:2303.15343*, 2023. 1
- [134] Renrui Zhang, Xiangfei Hu, Bohao Li, Siyuan Huang, Hanqiu Deng, Yu Qiao, Peng Gao, and Hongsheng Li. Prompt, generate, then cache: Cascade of foundation models makes strong few-shot learners. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 11
- [135] Yunhua Zhang, Hazel Doughty, and Cees GM Snoek. Low-resource vision challenges for foundation models. *arXiv preprint arXiv:2401.04716*, 2024. 11
- [136] George Kingsley Zipf. *Human behavior and the principle of least effort: An introduction to human ecology*. Ravenio books, 2016. 8

# Part I

## Appendix

### Table of Contents

---

<b>A</b>	<b>Concept Frequency is Predictive of Performance Across Prompting Strategies</b>	<b>2</b>
<b>B</b>	<b>Concept Frequency is Predictive of Performance Across Retrieval Metrics</b>	<b>3</b>
<b>C</b>	<b>Concept Frequency is Predictive of Performance for T2I Models</b>	<b>5</b>
<b>D</b>	<b>Concept Frequency is Predictive of Performance across Concepts only from Image and Text Domains</b>	<b>9</b>
<b>E</b>	<b>Experimental Details</b>	<b>11</b>
E.1	Setup of Mayilvahanan et al. [79] . . . . .	11
E.2	<i>Let It Wag!</i> : Test Set Curation . . . . .	11
<b>F</b>	<b>Why and How Do We Use RAM++?</b>	<b>12</b>
F.1	Why RAM++ and not CLIP or open-vocabulary detectors? . . . . .	12
F.2	How: Optimal RAM++ threshold for calculating concept frequencies . . . . .	12
<b>G</b>	<b>Details about Misalignment Degree Results</b>	<b>15</b>
<b>H</b>	<b>T2I Models: Evaluation</b>	<b>16</b>
H.1	Quantitative Results by Retrieval . . . . .	16
H.2	Qualitative Results . . . . .	17
<b>I</b>	<b>Classification Results: <i>Let It Wag!</i></b>	<b>22</b>

---

# A Concept Frequency is Predictive of Performance Across Prompting Strategies

We extend the zero-shot classification results from Fig. 2 in Fig. 8 with two different prompting strategies: the results in the main paper used the {classname} only as the prompts, here we showcase both (1) “A photo of a {classname}” prompting and (2) 80 prompt ensembles as used by Radford et al [91]. We observe that *the strong log-linear trend between concept frequency and zero-shot performance consistently holds across different prompting strategies.*

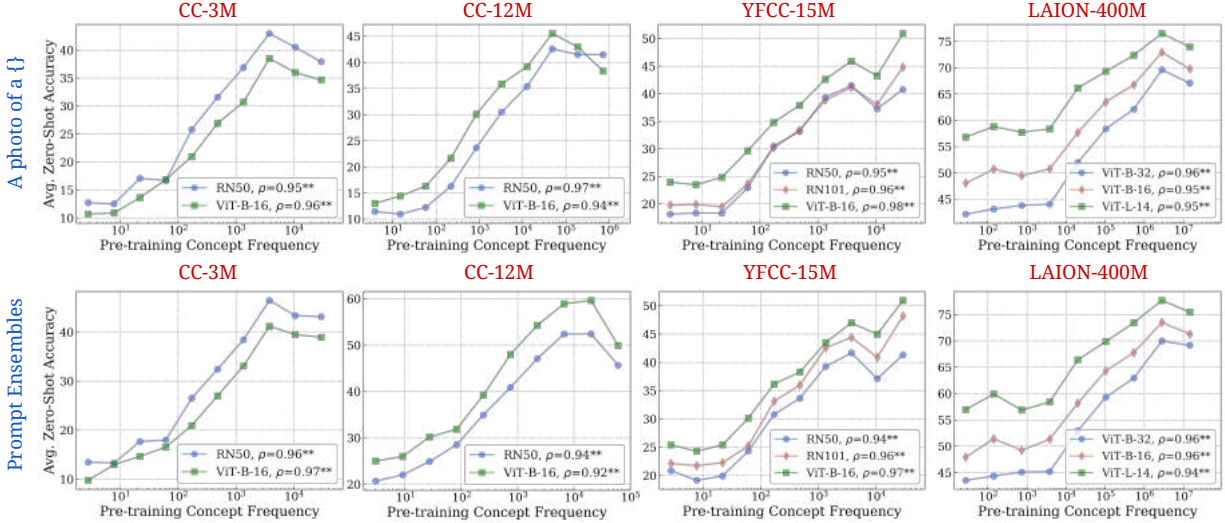


Figure 8: **Log-linear relationships between concept frequency and CLIP zero-shot performance.** Across all tested architectures (RN50, RN101, ViT-B-32, ViT-B-16, ViT-L-14) and pretraining datasets (CC-3M, CC-12M, YFCC-15M, LAION-400M), we observe a consistent linear relationship between CLIP’s zero-shot classification accuracy on a concept and the log-scaled concept pretraining frequency. This trend holds for both “A photo of a {classname}” prompting style and 80 prompt ensembles [91]. \*\* indicates that the result is significant ( $p < 0.05$  with a two-tailed t-test.), and thus we show pearson correlation ( $\rho$ ) as well.

## B Concept Frequency is Predictive of Performance Across Retrieval Metrics

We supplement Fig. 2 in the main paper, where we showed results with the text-to-image (I2T) recall@10 metric. In Figs. 9 and 10, we present results for the retrieval experiments across all six metrics: I2T-Recall@1, I2T-Recall@5, I2T-Recall@10, T2I-Recall@1, T2I-Recall@5, T2I-Recall@10. We observe that *the strong log-linear trend between concept frequency and zero-shot performance robustly holds across different retrieval metrics*.

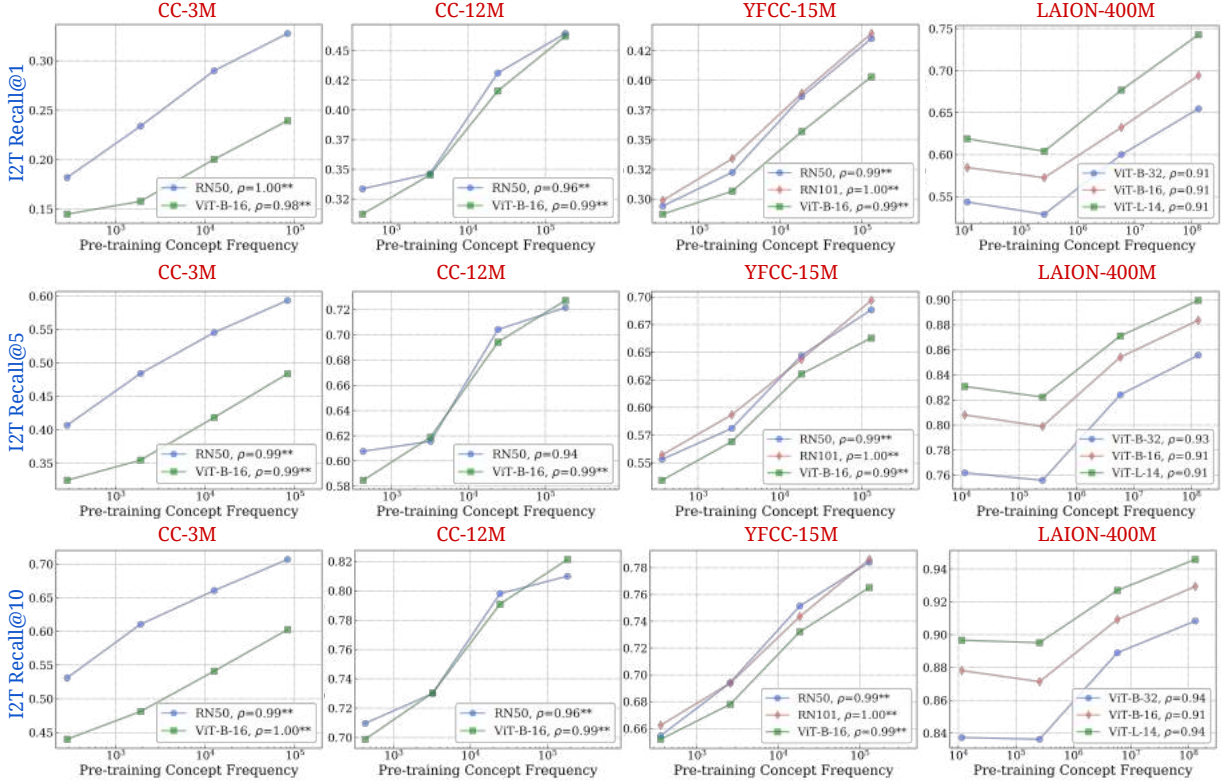


Figure 9: **Log-linear relationships between concept frequency and CLIP I2T retrieval performance.** Across all tested architectures (RN50, RN101, ViT-B-32, ViT-B-16, ViT-L-14) and pretraining datasets (CC-3M, CC-12M, YFCC-15M, LAION-400M), we observe a consistent linear relationship between CLIP’s retrieval performance (measured using image-to-text metrics) on a concept and the log-scaled concept pretraining frequency. \*\* indicates that the result is significant ( $p < 0.05$  with a two-tailed t-test.), and thus we show pearson correlation ( $\rho$ ) as well.

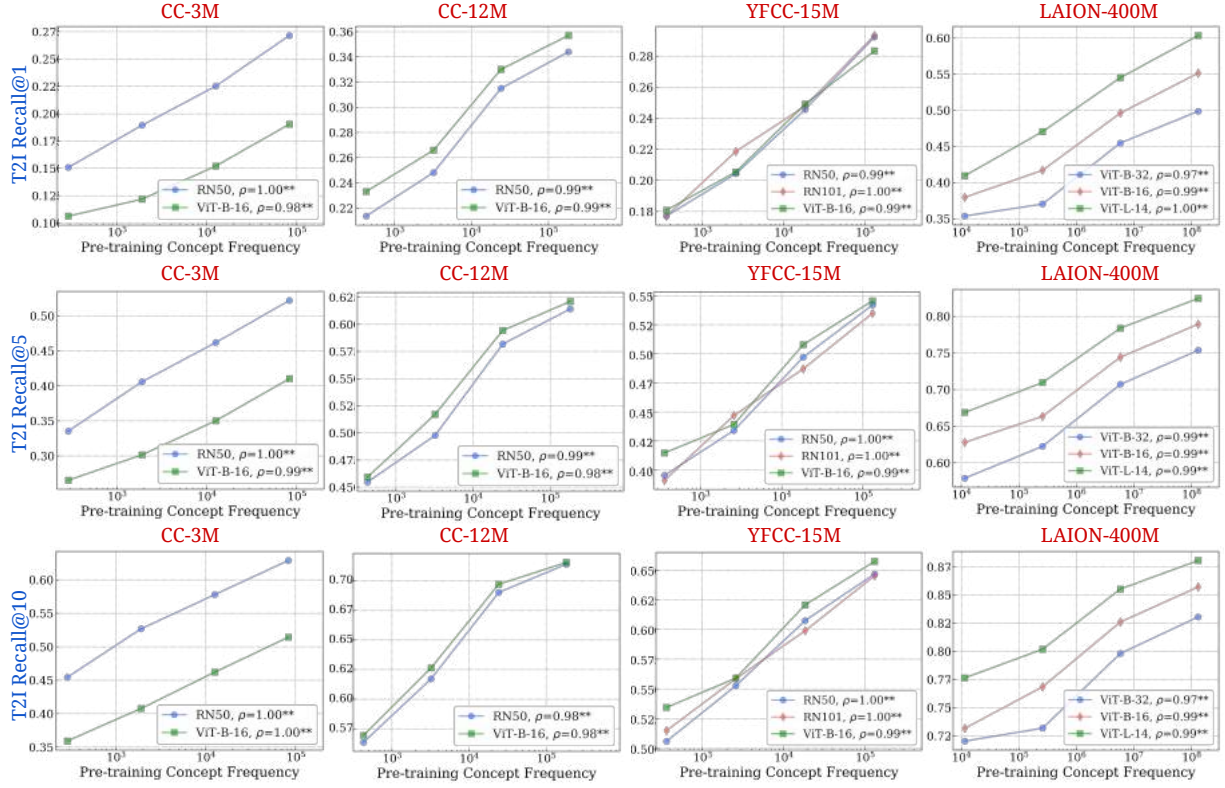


Figure 10: **Log-linear relationships between concept frequency and CLIP T2I retrieval performance.** Across all tested architectures (RN50, RN101, ViT-B-32, ViT-B-16, ViT-L-14) and pretraining datasets (CC-3M, CC-12M, YFCC-15M, LAION-400M), we observe a consistent linear relationship between CLIP’s retrieval performance (measured using text-to-image metrics) on a concept and the log-scaled concept pretraining frequency. \*\* indicates that the result is significant ( $p < 0.05$  with a two-tailed t-test.), and thus we show pearson correlation ( $\rho$ ) as well.



## C Concept Frequency is Predictive of Performance for T2I Models

We extend the results from Fig. 3 with Figs. 11 to 15. As with Fig. 3, due to the high concept frequency, the scaling trend is weaker. Furthermore, we do see inconsistency in the trends for the human-rated scores retrieved from HEIM [71], hence we perform a small scale human evaluation to check them.

Given the societal relevance [23], we decided to test Stable Diffusion [96] (v1.4) on generating public figures. We scraped 50,000 people from the “20230123-all” Wikidata JSON dump by filtering for entities listed as “human” [8], and scraped a reference image for the human study for each person if an image was available. After computing concept frequency from LAION-Aesthetics text captions (using suffix array [70]), we found that  $\approx 10,000$  people were present in the pretraining dataset. Note that to ensure the people’s names were treated as separate words, we computed frequency for strings of the format “{entity}”. We then randomly sample 360 people (for which a reference image was available) normalized by frequency [22] for the human study. For generating images with Stable Diffusion, we used the prompt “headshot of {entity}”, in order to specify to the model that “{entity}” is referring to the person named “{entity}” [50].

We assessed image-text alignment with a human study with 6 participants, where each participant was assigned 72 samples; for consistency, of the 360 total samples, we ensured 10% were assigned to 3 participants. Provided with a reference image, the participants were asked if the sample accurately depicts the prompt. Three choices were provided: “Yes” (score=1.), “Somewhat” (score=0.5), and “No” (score=0.). Accuracy was computed by averaging the scores.

As can be seen in Fig. 16, we observe a log-linear trend between concept frequency and zero-shot performance. Thus, we observe that *the log-linear trend between concept frequency and zero-shot performance consistently holds even for T2I models.*

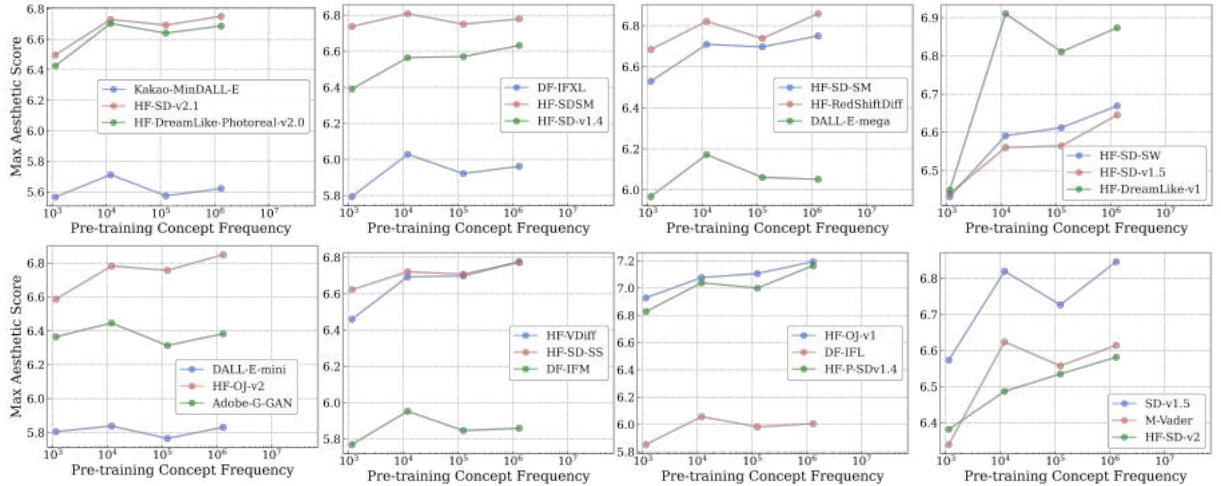


Figure 11: **Log-linear relationships between concept frequency and T2I Max aesthetic scores.** Across all tested models pretrained on the LAION-Aesthetics dataset, we observe a consistent linear relationship between T2I zero-shot performance on a concept and the log-scaled concept pretraining frequency.

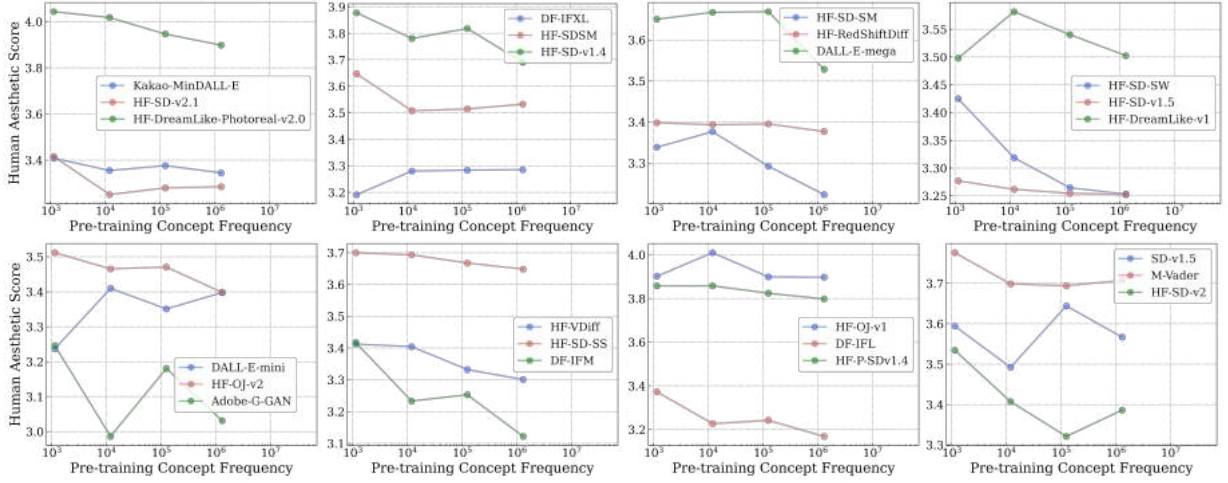


Figure 12: **Log-linear relationships between concept frequency and T2I human aesthetic scores.** Across all tested models pretrained on the LAION-Aesthetics dataset, we observe a consistent linear relationship between T2I zero-shot performance on a concept and the log-scaled concept pretraining frequency.

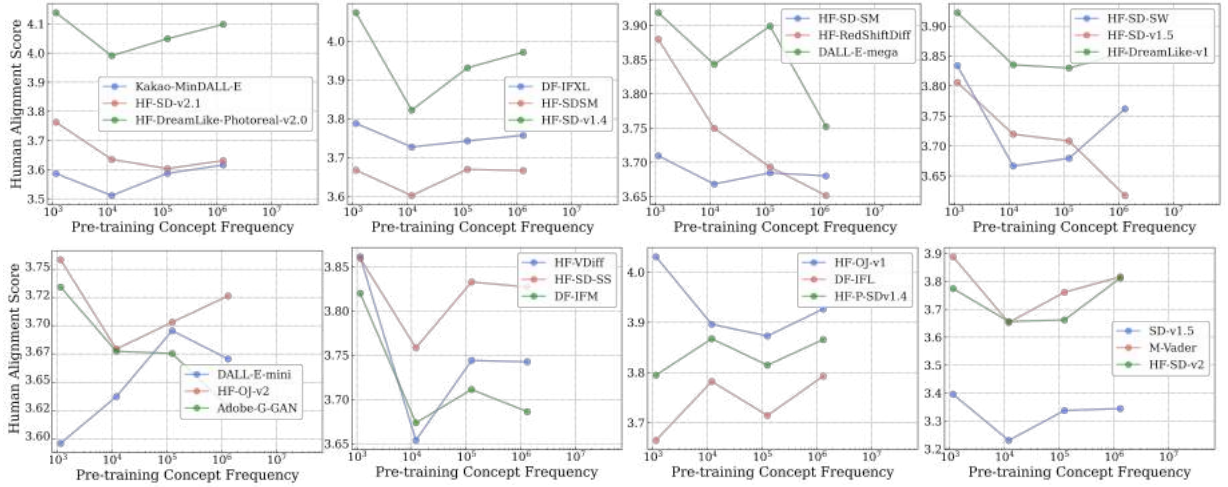


Figure 13: **Log-linear relationships between concept frequency and T2I human alignment scores.** Across all tested models pretrained on the LAION-Aesthetics dataset, we observe a consistent linear relationship between T2I zero-shot performance on a concept and the log-scaled concept pretraining frequency.

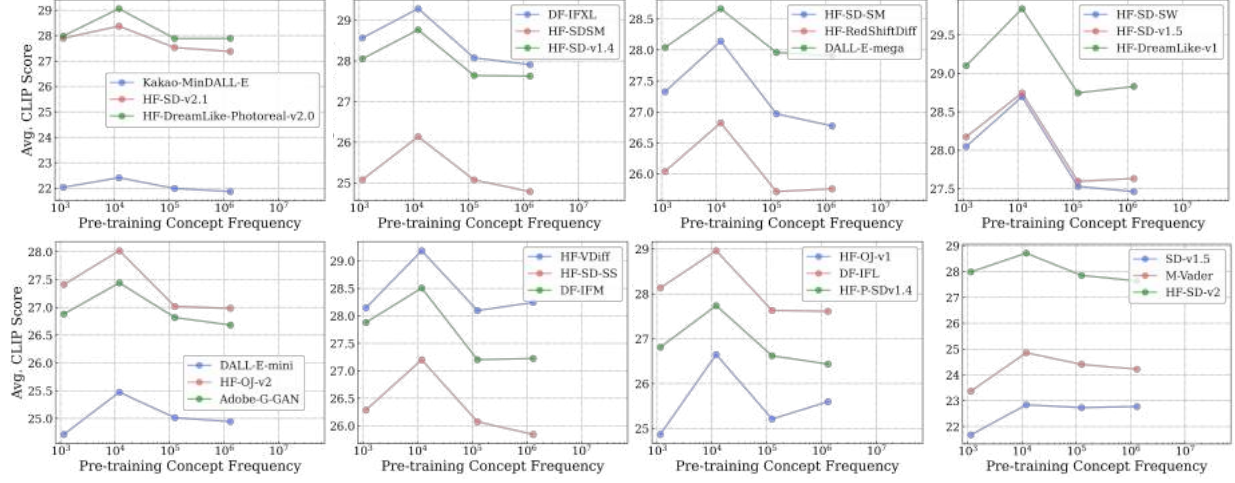


Figure 14: **Log-linear relationships between concept frequency and T2I Avg. CLIP scores.** Across all tested models pretrained on the LAION-Aesthetics dataset, we observe a consistent linear relationship between T2I zero-shot performance on a concept and the log-scaled concept pretraining frequency.

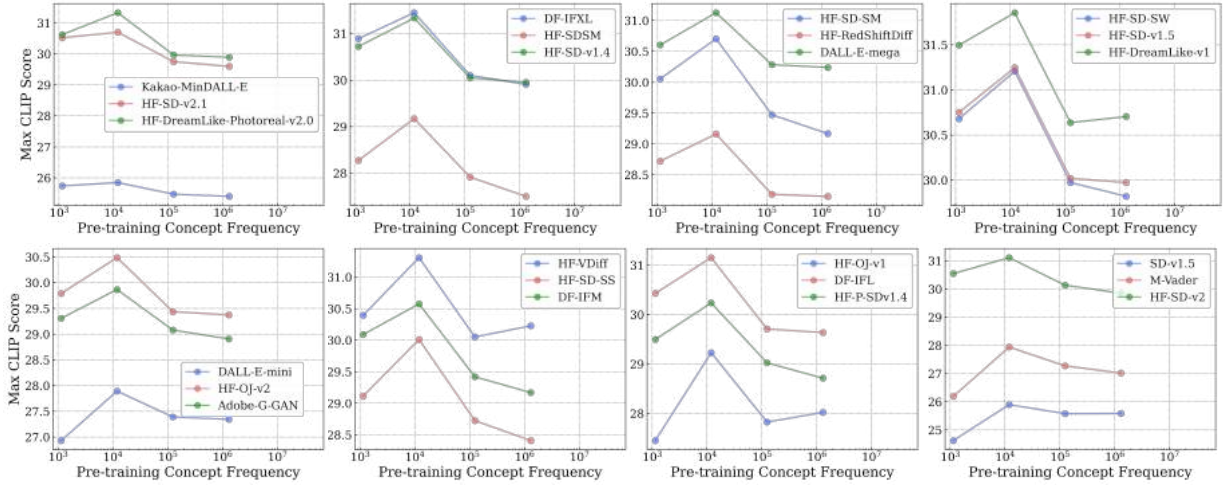


Figure 15: **Log-linear relationships between concept frequency and T2I Max CLIP scores.** Across all tested models pretrained on the LAION-Aesthetics dataset, we observe a consistent linear relationship between T2I zero-shot performance on a concept and the log-scaled concept pretraining frequency.




Figure 16: **Log-linear relationship between concept frequency and T2I human evaluation for text-image alignment for people concepts.** We observe a consistent linear relationship between T2I zero-shot performance on a concept and the log-scaled concept pretraining frequency.

## D Concept Frequency is Predictive of Performance across Concepts only from Image and Text Domains

In all the main performance-frequency plots we have presented until now, the concept frequencies were estimated using the intersection of the image-frequencies and the text-frequencies. Here, we showcase results with using them independently in Figs. 17 and 18 respectively. We note that both independent searching methods showcase log-linear trends as before confirming our main result. We observe that *the strong log-linear trend between concept frequency and zero-shot performance robustly holds across concepts derived from image and text domains independently as well.*

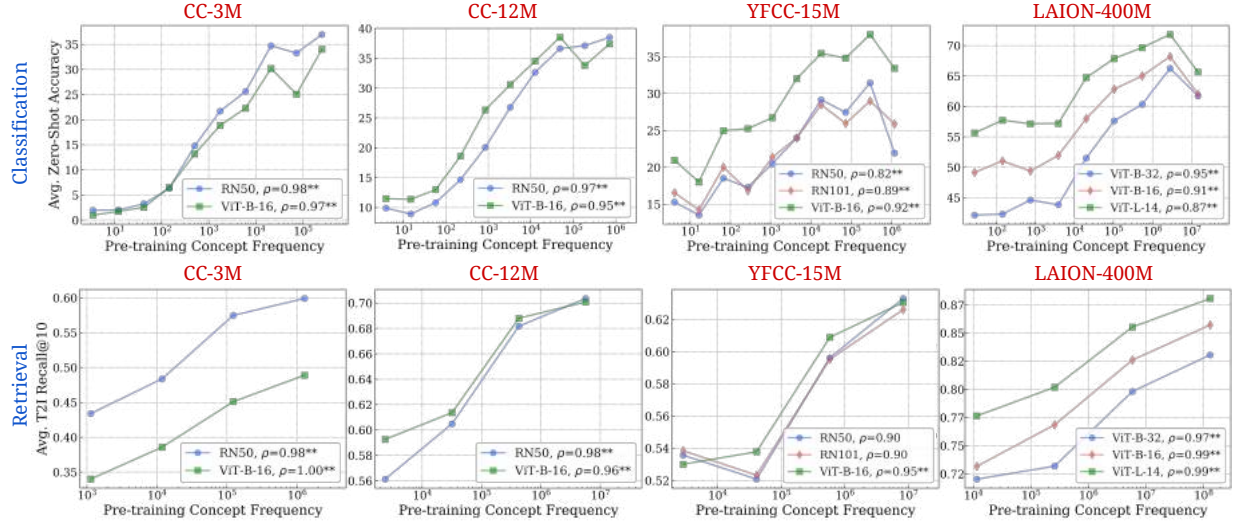


Figure 17: **Log-linear relationships between image concept frequency and CLIP performance.** Across all tested architectures (RN50, RN101, ViT-B-32, ViT-B-16, ViT-L-14) and pretraining datasets (CC-3M, CC-12M, YFCC-15M, LAION-400M), we observe a consistent linear relationship between CLIP’s zero-shot accuracy and retrieval performance on a concept and the log-scaled concept pretraining frequency (searched using only pretraining images). \*\* indicates that the result is significant ( $p < 0.05$  with a two-tailed t-test.), and thus we show pearson correlation ( $\rho$ ) as well.



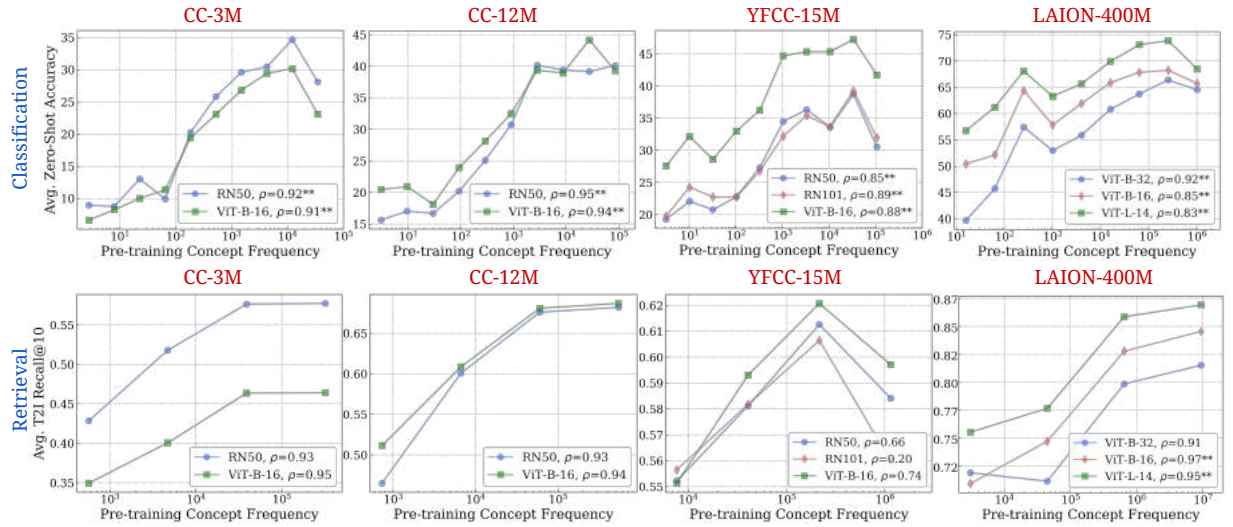


Figure 18: **Log-linear relationships between text concept frequency and CLIP performance.** Across all tested architectures (RN50, RN101, ViT-B-32, ViT-B-16, ViT-L-14) and pretraining datasets (CC-3M, CC-12M, YFCC-15M, LAION-400M), we observe a consistent linear relationship between CLIP’s zero-shot accuracy and retrieval performance on a concept and the log-scaled concept pretraining frequency (searched using only pretraining text captions). **\*\*** indicates that the result is significant ( $p < 0.05$  with a two-tailed t-test.), and thus we show pearson correlation ( $\rho$ ) as well.

## E Experimental Details

### E.1 Setup of Mayilvahanan et al. [79]

LAION-200M is a dataset obtained by deduplicating LAION-400M by pruning exact duplicates, near duplicates, and semantically similar samples within LAION-400M [10]. The control pretraining set is created by pruning 50 million highly similar samples from LAION in the order of decreasing perceptual similarity to datapoints in ImageNet-val set. We use the 150M pretraining set for obtaining the concept distribution. We evaluate the performance of a ViT-B/32 CLIP model trained on this dataset on our downstream tasks, and present our analysis on those tasks.

### E.2 *Let It Wag!*: Test Set Curation

To ensure our datasets are thoroughly cleaned and diverse, we follow a meticulous process:

**1. Diverse Sourcing:** We gather images from three different online sources—Flickr, DuckDuckGo, and Bing Search—to maximize the variety of our dataset, while retaining very easy-to-classify images<sup>2</sup>.

**2. Temporal Filtering:** We applied a filter to only retrieve images after January 2023 to minimize overlap with images used in the pre-training of Vision-Language Models (VLMs). Note this helps mitigate but does not ensure the overlap problem is resolved.

**3. Outlier Removal:** We employ a pre-trained InceptionNet [111] to remove outliers from the entire image pool. We do this by taking all pairwise cosine-similarities between all images in the pool, and removing the images that are in the bottom 5% of the similarity values<sup>3</sup>.

**4. Initial De-duplication with an InceptionNet:** We employ a pre-trained InceptionNet [111] model to identify and remove duplicates. This step involves setting high thresholds for soft de-duplication (0.9 for common classes and 0.95 for fine-grained classes) to ensure only minor, precise exclusions. A threshold of 0.9/0.95 means that we consider images to be duplicates if the cosine similarity of that image’s embedding (from InceptionNet) with any other image’s embedding in the image pool is larger than 0.9/0.95.

**5. Manual Verification:** Following the automated cleaning, we manually inspect and verify the accuracy of the remaining images for each class to ensure they meet quality standards.

**6. Second-level De-duplication with Perceptual Hashing:** Post-verification, we use perceptual hashing [37] with a threshold of 10 bits to identify and remove duplicate images within each class, ensuring uniqueness across our dataset<sup>4</sup>.

**7. Class Balancing:** Finally, we balance the dataset to ensure an equal representation of classes.

This process was followed for increased quality and reliability of our dataset for image recognition tasks.

---

<sup>2</sup>we use the image sourcing pipeline of C2C [90]

<sup>3</sup>We use the `fastdup` library for outlier removal: <https://github.com/visual-layer/fastdup>

<sup>4</sup>We use the `imagededup` library for de-duplication: <https://github.com/idealo/imagededup>

## F Why and How Do We Use RAM++?

We detail why we use the RAM++ model [59] instead of CLIPScore [56] or open-vocabulary detection models [80]. Furthermore, we elaborate on how we selected the threshold hyperparameter used for identifying concepts in images.

### F.1 Why RAM++ and not CLIP or open-vocabulary detectors?

We provide some qualitative examples to illustrate why we chose RAM++. Our input images do not often involve complex scenes suitable for object detectors, but many fine-grained classes on which alongside CLIP, even powerful open-world detectors like OWL-v2 [80] have poor performance.



Figure 19: **Qualitative Results comparing OWL-v2, RAM++ and CLIP.** We show qualitative examples across three different models: OWL-v2, RAM++ and CLIP on fine-grained concepts.

### F.2 How: Optimal RAM++ threshold for calculating concept frequencies

We ablate the choice of the threshold we use for assigning concepts to images using the RAM++ model. For the given set of concepts, RAM++ provides a probability value (by taking a sigmoid over raw logits) for each concept’s existence in a particular image. To tag an image as containing a particular concept, we have to set a threshold deciding this assignment. We test over three thresholds: {0.5, 0.6, 0.7}, showcasing quantitative and qualitative results for all thresholds in Figs. 20 and 21.

We observe best frequency estimation results using the highest frequency of 0.7. This is due to the high precision afforded by this threshold, leading to us counting only the “most aligned images” per concept as hits. With lower thresholds (0.5, 0.6), we note that noisier images that do not align well with the concept can be counted as hits, leading to degraded precision and thereby poorer frequency estimation. Hence, we use 0.7 as the threshold for all our main results.

RAM++ Threshold			
			
0.5	<p>school bus   trolleybus   stands   Highway or Road   alley   courtyard   crosswalk   driveway   park   parking_lot   plaza   street   license-plate   people   school-bus   car   truck   bus   man   road   floor   stop   bunch   step   crowd   couple   corner   travel   white   pair   outside   male   size   traffic   time   vehicle   turn   up   path   square   van   guy   close   photo   store   middle   roof   lot   place   ride   area   men   rain   red   sidewalk   number   pavement   walkway   air   intersection   direction   building   market   ground   image   day   walk   cross   town   passenger   side   commuter   setting   move   color   blanket   stand   water   green   picture   curb   drive   figure   kind   pedestrian   way   city   line   person   roadway</p>	<p>baluster / handrail   spiral or coil   stone wall   couch   television   window screen   studio   lights   spiralled   dining_room   house   kitchen   living_room   ...   plant   brown   wall   fireplace   window   room   railing   rail   show   light   furniture   image   side   fire   design   color   ceiling   brick   pict\$   re   figure   door   style   sculpture   display</p>	<p>lighthouse   dock   fountain   greenhouse   patio   pedestal   pier   pole   umbrella   w   ebsite   circle   her   being   stands   diagram   length   shape   illustration   drawing   Forest   Herbaceous Vegetation Land   River   Sea or Lake   boardwalk   ...   landscape   walk   cross   tree   part   life   surface   structure   space   side   garden   pool   land   design   setting   color   item   blanket   stand   water   green   reflection   model   picture   bench   figure   kind   way   city   line   work   sculpture   body   display   person   field   planter   living   spot</p>
0.6	<p>Highway or Road   park   parking_lot   plaza   street   car   truck   bus   man   road   stop   crowd   couple   travel   white   outside   traffic   vehicle   ride   rain   red   sidewalk   number   pavement   building   walk   cross   town   passenger   side   drive   pedestrian   way   city   line   roadway</p>	<p>baluster / handrail   lights   house   living_room   staircase   flooring   floor   step   frame   white   home   dark   stair   apartment   wall   fireplace   room   railing   1   ight   furniture   image   fire   ceiling   brick   picture</p>	<p>dock   fountain   pier   website   circle   stands   diagram   illustration   River   Sea or Lake   courtyard   park   pond   river   people   man   sea   floor   step   couple   white   island   outside   grass   bush   path   lake   place   inside   area   plant   variety   walkway   flower   show   building   piece   ground   image   landscape   tree   structure   side   garden   design   setting   stand   water   green   picture   bench   figure   way   body   person</p>
0.7	<p>park   parking_lot   street   bus   road   traffic   red   pedestrian   city   roadway</p>	<p>house   living_room   floor   step   white   home   stair   wall   fireplace   room   light</p>	<p>park   pond   couple   island   outside   plant   flower   show   image   tree   garden   design   water   green   picture</p>

Figure 20: **Qualitative Results with different RAM++ thresholds.** We show qualitative examples across three different thresholds:  $\{0.5, 0.6, 0.7\}$  for estimating concept frequency using the RAM++ model. We note that the significantly better concepts identified by the higher threshold (0.7) compared to the lower thresholds (0.5, 0.7). The images are sourced from the CC-3M dataset.

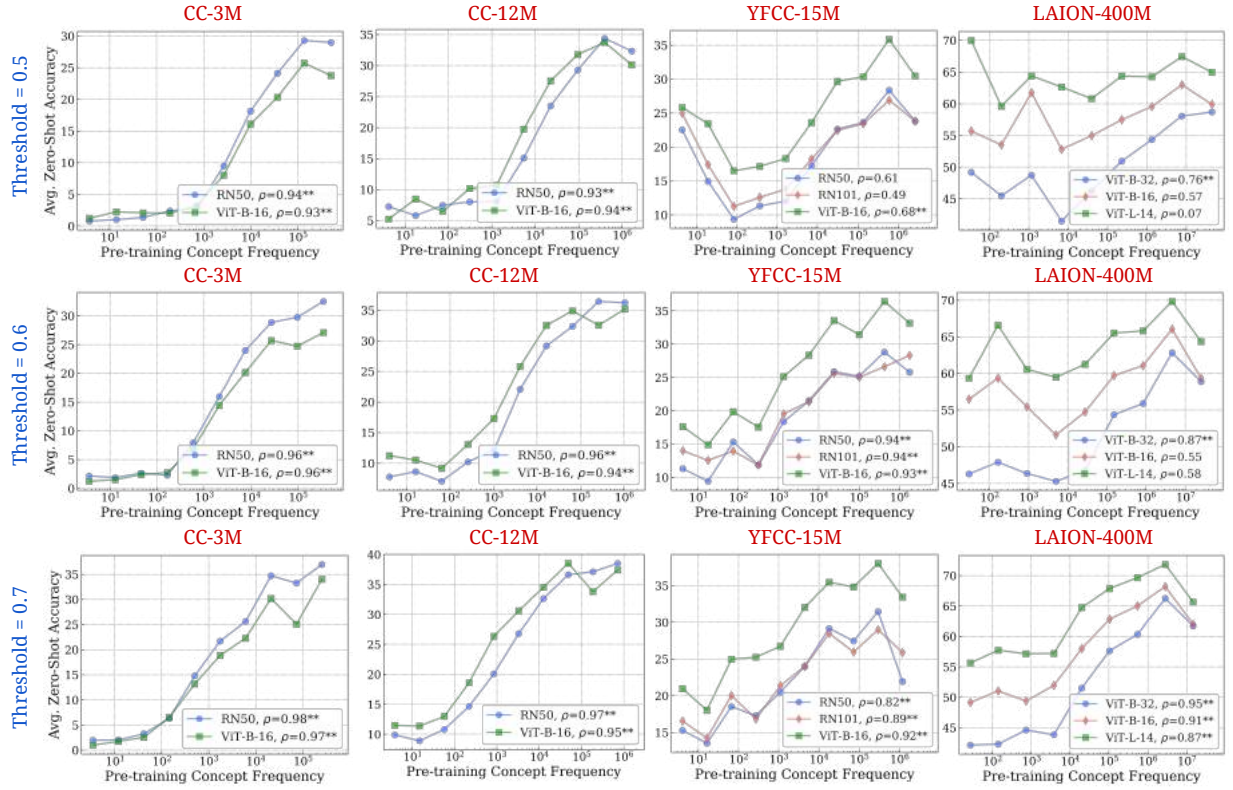


Figure 21: **Effect of different thresholds for determining concept frequency using RAM++.** We test three different thresholds:  $\{0.5, 0.6, 0.7\}$  for estimating concept frequency using the RAM++ model. We note that the correlations are significantly stronger with a threshold of 0.7—this is justified by the higher precision of image sample hits at a higher threshold (0.7). Comparatively, lower thresholds (0.5, 0.7) lead to noisier images being counted as hits, hence reducing the hit precision for determining frequency. \*\* indicates that the result is significant ( $p < 0.05$  with two-tailed t-test.), and thus we show pearson correlation ( $\rho$ ) too.







## G Details about Misalignment Degree Results

In Tab. 3 in the main paper, we quantified the *misalignment degree*, and showcased that a large number of image-text pairs in all pretraining datasets are misaligned. In Alg. 1, we describe the method used for quantifying the *misalignment degree* for each pretraining dataset. We also showcase some qualitative examples of a few image-text pairs from the CC-3M dataset that are identified as misaligned using our analysis.

<b>Data:</b> Pretraining dataset $\mathcal{D} = \{(i_1, t_1), (i_2, t_2), \dots, (i_N, t_N)\}$ , Image Index $I_{\text{img}}$ , Text Index $I_{\text{text}}$
<b>Result:</b> <i>mis_degree</i>
<pre> mis_degree ← 0 for (i, t) ∈ D do     img_concepts ← I_img[i] // extract all concepts from this image     text_concepts ← I_text[t] // extract all concepts from this text caption     hits ← set_intersection(img_concepts, text_concepts)     if len(hits) = 0 then         mis_degree ← mis_degree + 1     end end return mis_degree/N end </pre>

**Algorithm 1:** Extracting *misalignment degree* from pretraining datasets

Identified misaligned image-text pair				
	<b>Caption:</b> According to the model, she regularly gets told it looks	<b>Caption:</b> person, why u turn me on with your designs? i believe i will begining some more of his work here in a bit.	<b>Caption:</b> put all the extraordinary things that are not ordinary, simple things to do everything on that simple.	<b>Caption:</b> step by step ... 2
Concepts from image-search	['body', 'side', 'arm', 'female', 'woman', 'shoulder', 'back', 'jacket', 'dress', 'blue', 'male', 'neck']	['eye', 'art', 'monster', 'poster', 'red', 'green', 'gasmask']	['water_lilly', 'water lily', 'leave', 'pond', 'lotus', 'plant', 'buttercup', 'flower', 'water', 'green', 'closeup', 'white']	['head', 'eye', 'picture', 'shape', 'sketch', 'female', 'drawing', 'girl', 'smile', 'woman', 'faces-easy', 'face']
Concepts from text-search	['model']	['person', 'work', 'design']	['thing']	['step']

**Figure 22: Qualitative examples of misaligned image-text pairs identified.** We present 4 samples from the CC3M pretraining dataset that are identified as misaligned by our analysis. Here, the text captions clearly do not entail the images, and hence do not provide a meaningful signal for learning.

## H T2I Models: Evaluation

We provide additional quantitative and qualitative results in this section for T2I models evaluated on the “*Let It Wag!*” dataset.

### H.1 Quantitative Results by Retrieval

We provide further analysis on how state-of-the-art T2I models perform on the long-tailed concepts comprising the “*Let It Wag!*” dataset. As detailed in Sec. 6, we generate 4 images for each concept using Stable Diffusion XL [89], Stable Diffusion v2 [96] and Dreamlike Photoreal [1].

**Prompting Strategy.** The prompting strategy (system role) used, adapted from Shahmohammadi et al. [105], was:

Follow my commands:

1. I wish to generate text prompts about a given subject which I will use for image generation using off-the-shelf text-to-image models such as Stable Diffusion and DALL-E 3.
2. Assume all the subjects are nouns.
3. Follow a similar style and length of prompts as coco-captions.
4. Keep prompts concise and avoid creating prompts longer than 40 words.
5. Structure all prompts by setting a scene with at least one subject and a concrete action term, followed by a comma, and then describing the scene. For instance, “a view of a forest from a window in a cosy room, leaves are falling from the trees.”

Generate detailed prompts for the concepts in the order in which they are given. Your output should be just the prompts, starting with “1.”

With this pool of generated images, we conduct a controlled experiment on the long-tailed concepts using nearest-neighbor retrieval as the evaluation metric by querying a generated image and retrieving the top-k results from a gallery of images taken from the “*Let It Wag!*” dataset. The overall pipeline is as follows:

**Setup.** We define the query and gallery set for head and tail concepts. For tail concepts, we sample the 25 concepts with the lowest frequency from the “*Let It Wag!*” dataset. For head concepts, we sample the 25 most frequent concepts for comparison. We use the same prompting strategy with the selected 25 concepts across all 3 T2I models. To create the gallery set, we randomly sample 100 images for each of these concepts. We use DINOv2 [85] ViT-S/14 as the feature extractor.

**Results.** In Table 5, we provide the Cumulative Matching Characteristic (CMC@k) results for all 3 T2I models used in our experiment. CMC@k was chosen as we are interested in measuring the delta between head and tail concepts for successful retrievals within the top-k retrieved real images for a given generated image. We observe a large performance gap between *Head* and *Tail* concepts, providing a quantitative evaluation of generation performance of T2I models.

Table 5: **Generated-real retrieval scores.** We compare retrieval results of DINOv2 ViT-S/14 when using generated images as query images. We report  $\Delta$  CMC@k results where  $k=\{1,2,5\}$  between head and tail concepts.

Model	$\Delta$ CMC		
	k=1	k=2	k=5
Stable Diffusion XL	13.0	16.0	16.8
Stable Diffusion v2	11.0	10.0	10.4
Dreamlike Photoreal	8.0	9.0	9.4

## H.2 Qualitative Results

In Fig. 7 of the main text, we provide an initial insight into the qualitative performance of T2I models on “*Let It Wag!*” concepts. For ease of comprehension and comparison, we segregate concepts into 4 clusters: Aircraft (Fig. 23), Activity (Fig. 24), Animal (Fig. 25) and others (Fig. 26).

**Results.** Fig. 23 shows T2I models having difficulty in representing an aircraft in its full form in a majority of cases in addition to misrepresenting the specific model in the generated images. Fig. 24 showcases the difficulty T2I models face when representing actions or activities from prompts. Fig. 25 exemplifies the same inability of T2I models to accurately represent animal species. Finally, the remainder of the query set is shown in Fig. 26 and includes the inability to classify and subsequently generate certain species of flowers and objects.

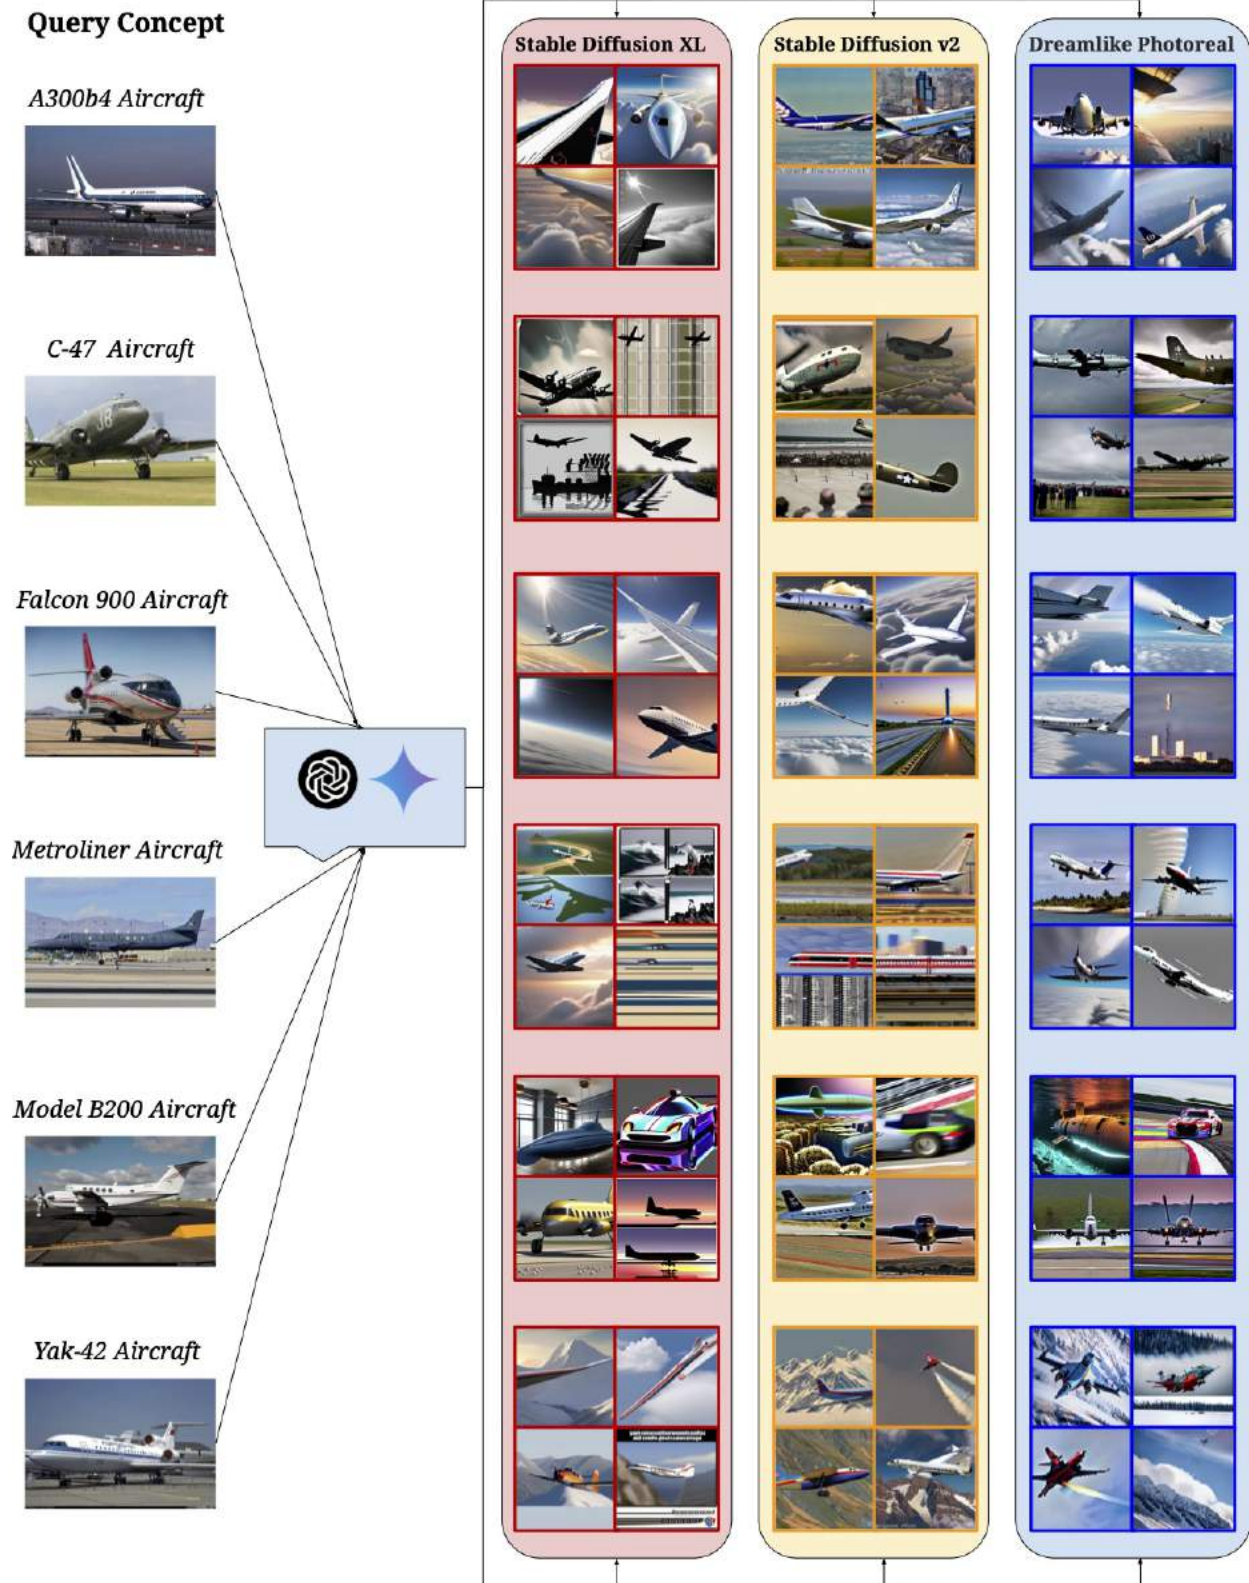


Figure 23: Qualitative results on the Aircraft cluster.



## Query Concept

*Cricket Bowling*



*Head Massage*



*Juggling Balls*



*Playing Daf*



*Soccer Juggling*



*Wall Pushups*

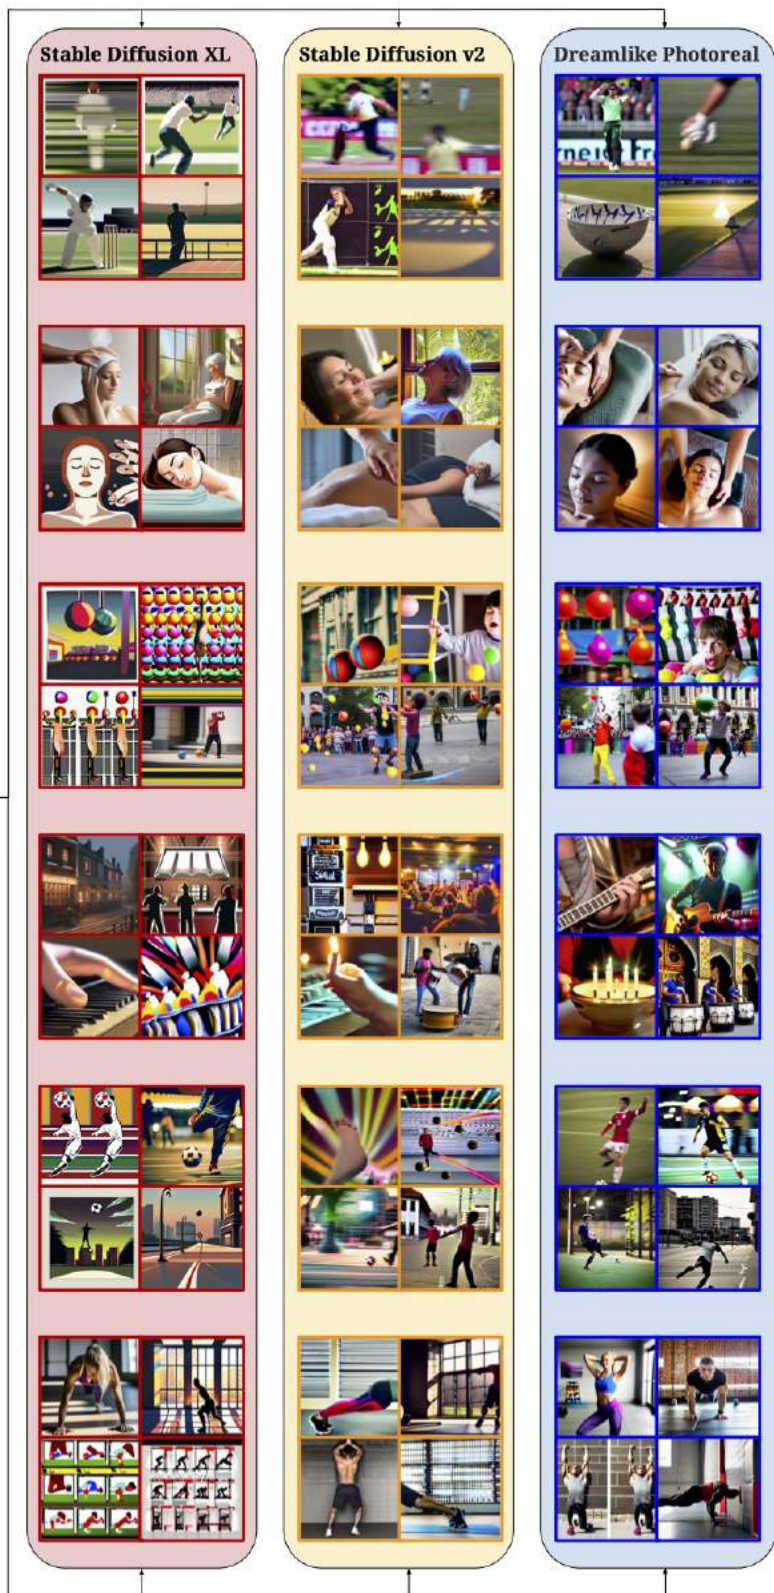


Figure 24: Qualitative results on the Activity cluster.



## Query Concept

*Cave Swallow Bird*



*Clark's Grebe Bird*



*Common Grackle Bird*



*Crested Auklet Bird*



*Chiton*



*Eel*

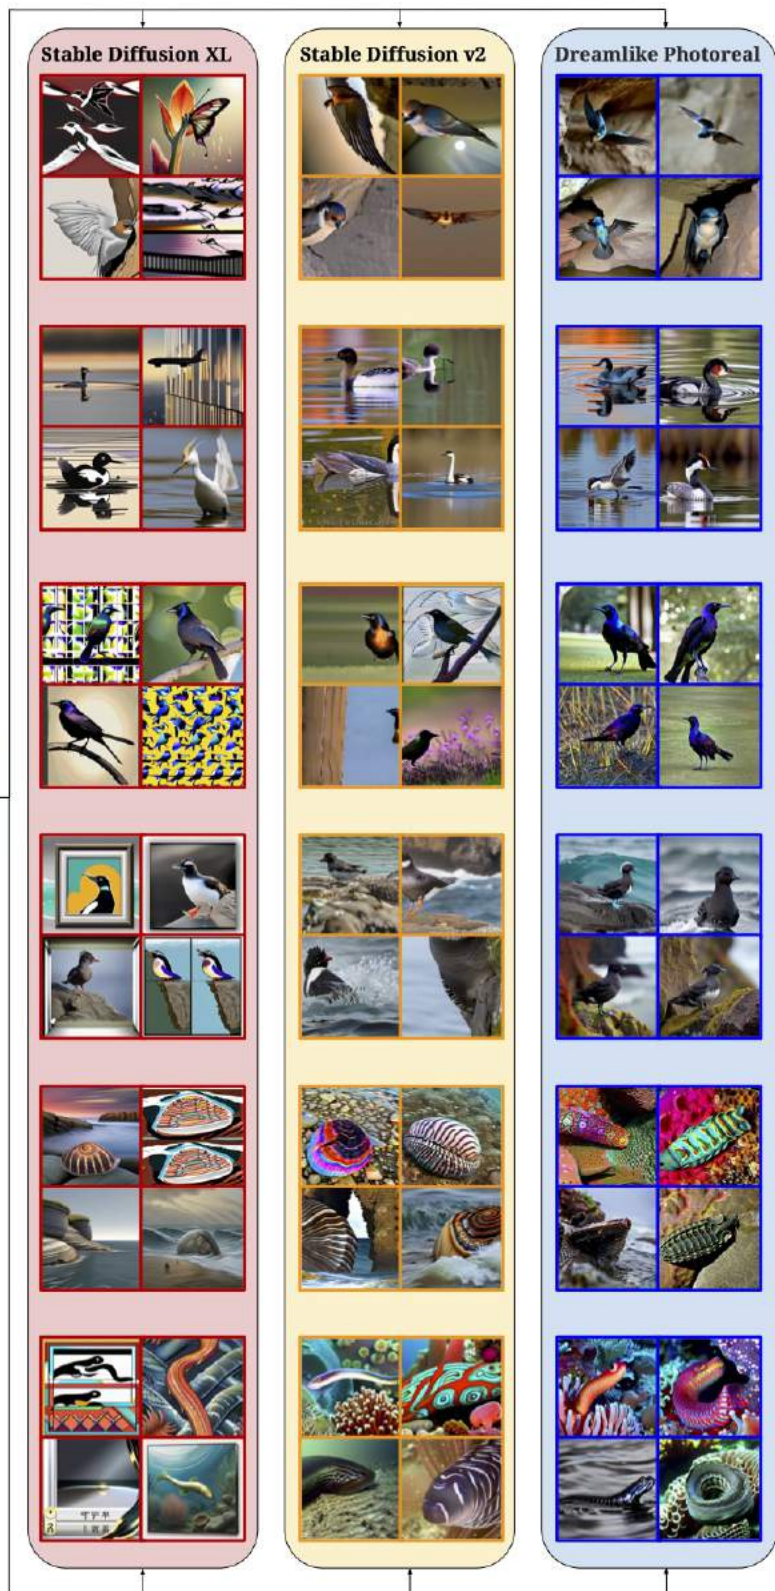


Figure 25: Qualitative results on the Animal cluster.

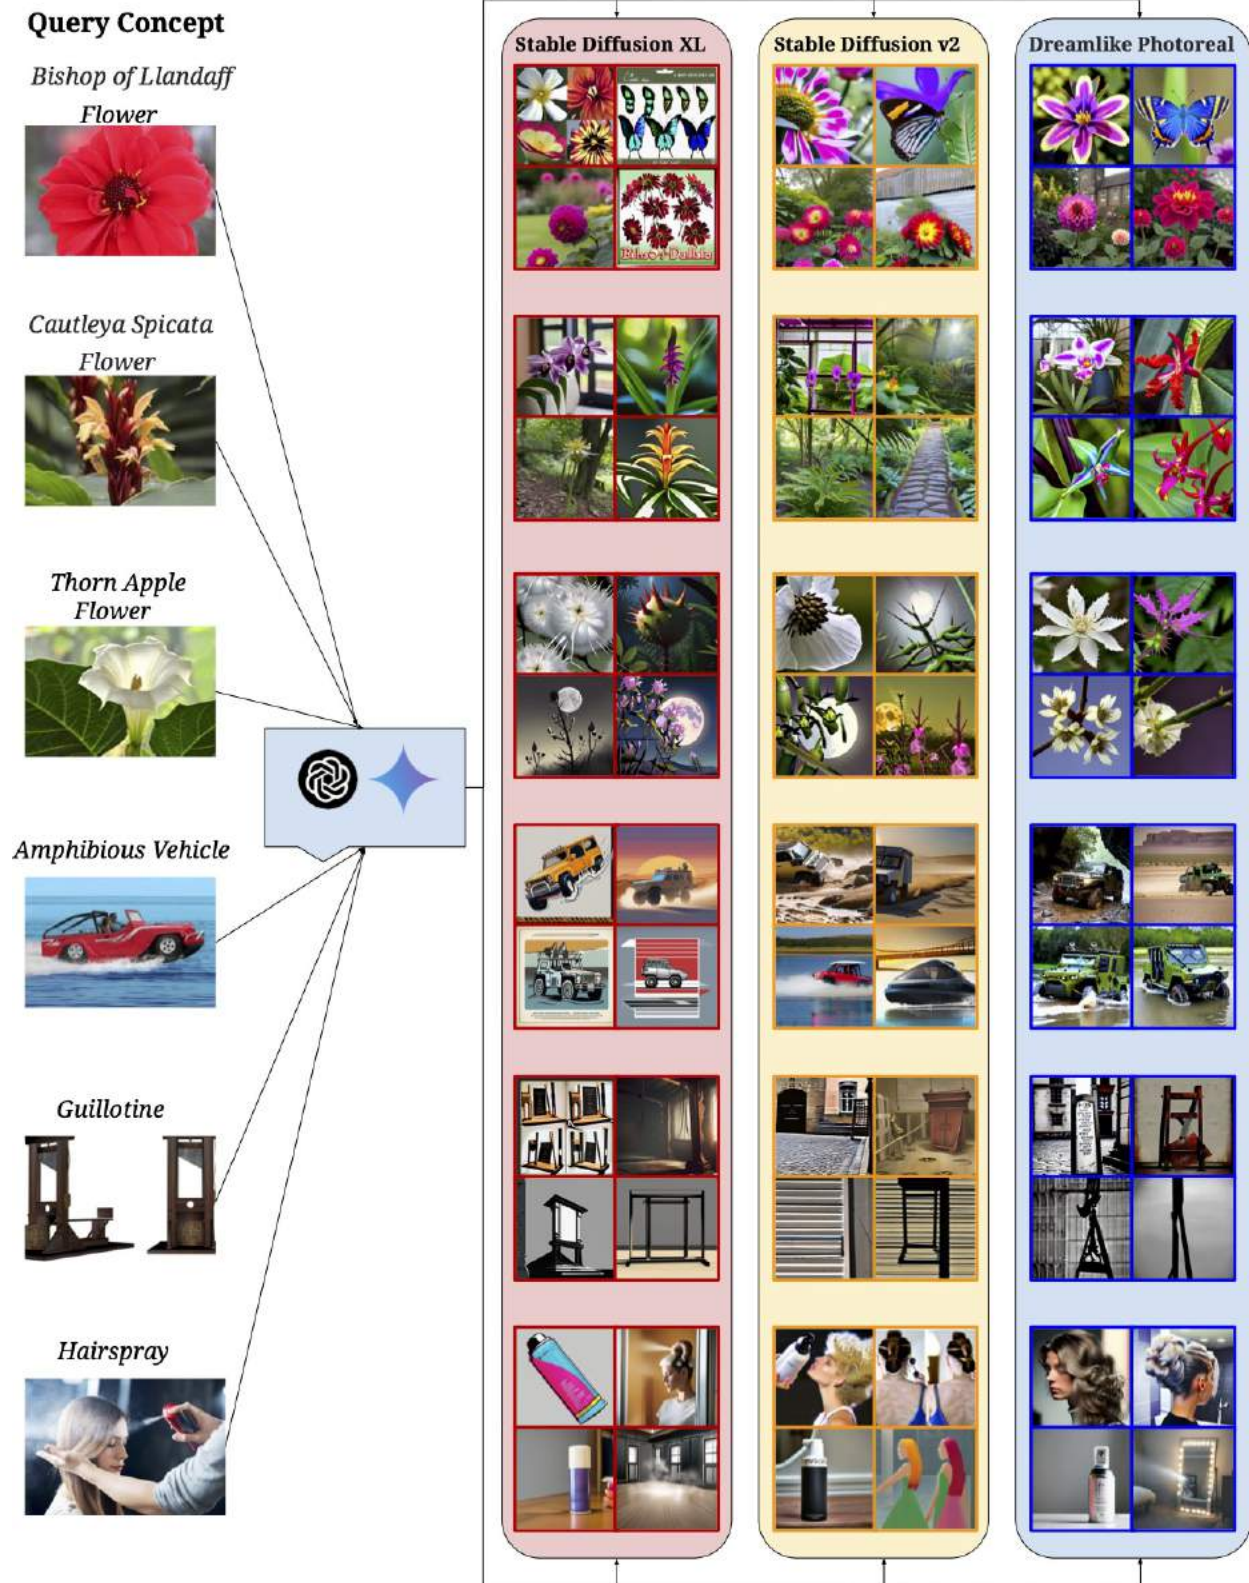


Figure 26: Qualitative results for other selected failure cases.



# I Classification Results: *Let It Wag!*

Here, we present the raw accuracy values of the 40 tested models on both *Let It Wag!* and ImageNet in Tab. 6. For reference, we also report the datasets these models were trained on and the number of parameters for each model. We see clear drops in performance compared to ImageNet, across model sizes, architectures and pretraining datasets.

Table 6: Full results dump on *Let It Wag!* and ImageNet.

Pretraining Dataset	Model	Num. Parameters (in millions)	ImageNet Acc.	<i>Let It Wag!</i> Acc.
CC-3M [107]	RN50	102.01	20.09	3.74
	ViT-B-16	149.62	17.10	3.01
CC-12M [27]	RN50	102.01	33.14	8.92
	ViT-B-16	149.62	37.39	11.49
YFCC-15M [113]	RN50	102.01	31.88	13.15
	RN101	119.69	34.04	15.19
	ViT-B-16	149.62	37.88	19.25
OpenAI-WIT [91]	RN50	102.01	59.82	31.93
	RN101	119.69	62.28	31.88
	ViT-B-32	151.28	63.32	33.52
	ViT-B-16	149.62	68.34	37.85
	ViT-L-14	427.62	75.54	45.31
WebLI [29]	ViT-B-16	203.79	78.49	54.63
	ViT-L-16	652.15	82.07	61.50
	SO400M	877.36	83.44	67.32
DataComp [46]	ViT-B-32	151.28	69.18	46.90
	ViT-B-16	149.62	73.48	52.89
	ViT-L-14	427.62	79.21	63.04
DataComp-DFN [43]	ViT-B-16	149.62	76.24	56.59
	ViT-H-14	986.11	83.44	71.91
CommonPool [46]	ViT-B-32	151.28	23.04	7.73
	ViT-B-16	149.62	57.77	20.97
	ViT-L-14	427.62	76.37	46.96
LAION-400M [102]	ViT-B-32	151.28	60.23	32.88
	ViT-B-16	149.62	67.02	39.13
	ViT-L-14	427.62	72.74	46.59
LAION-2B [103]	ViT-B-32	151.28	66.55	41.79
	ViT-B-16	149.62	70.22	44.21
	ViT-L-14	427.62	75.25	51.03
	ViT-H-14	986.11	77.92	58.98
	ViT-g-14	1366.68	78.46	59.01
	ViT-bigG-14	2539.57	80.09	63.54
MetaCLIP-400M [125]	ViT-B-32	151.28	65.58	40.50
	ViT-B-16	149.62	70.80	46.50
	ViT-L-14	427.62	76.20	52.78
MetaCLIP-FullCC [125]	ViT-B-32	151.28	67.66	43.84
	ViT-B-16	149.62	72.12	49.32
	ViT-L-14	427.62	79.17	57.48
	ViT-H-14	986.11	80.51	62.59
SynthCI-30M [51]	ViT-B-16	149.62	30.67	9.15