


# Dataset Growth

Ziheng Qin<sup>\*1</sup> , Zhaopan Xu<sup>\*1,2</sup>, Yukun Zhou<sup>1</sup>, Zangwei Zheng<sup>1</sup>, Zebang Cheng<sup>3</sup>, Hao Tang<sup>4</sup>, Lei Shang<sup>5</sup>, Baigui Sun<sup>5</sup>, Xiaojiang Peng<sup>3</sup>, Radu Timofte<sup>4</sup>, Hongxun Yao<sup>†2</sup>, Kai Wang<sup>†1</sup>, and Yang You<sup>†1</sup>

Code: <https://github.com/NUS-HPC-AI-Lab/InfoGrowth>

<sup>1</sup> National University of Singapore <sup>2</sup> Harbin Institute of Technology

<sup>3</sup> Shenzhen Technology University <sup>4</sup> ETH Zurich <sup>5</sup> Alibaba Group

**Abstract.** Deep learning benefits from the growing abundance of available data. Meanwhile, efficiently dealing with the growing data scale has become a challenge. Data publicly available are from different sources with various qualities, and it is impractical to do manual cleaning against noise and redundancy given today’s data scale. There are existing techniques for cleaning/selecting the collected data. However, these methods are mainly proposed for offline settings that target one of the cleanliness and redundancy problems. In practice, data are growing exponentially with both problems. This leads to repeated data curation with sub-optimal efficiency. To tackle this challenge, we propose InfoGrowth, an efficient online algorithm for data cleaning and selection, resulting in a growing dataset that keeps up to date with awareness of cleanliness and diversity. InfoGrowth can improve data quality/efficiency on both single-modal and multi-modal tasks, with an efficient and scalable design. Its framework makes it practical for real-world data engines.

## 1 Introduction

Deep learning benefits from the growing abundance of available data. Various datasets [1, 4, 12, 13, 26, 28] have been created and shared, contributing to the development of deep learning applications. In general, the scale of available data keeps growing (Fig. 1a web data growth). Even for a single task, there would be multiple datasets with a growing size over time. It has been almost impractical to manually maintain nowadays datasets. How to efficiently deal with the growing scale of data is becoming a challenge.

Specifically, many prominent datasets [4, 18] were collected from the Internet and then cleaned up by human annotators at a high cost. However, data continues to emerge on the Internet. According to [6], 0.33 zettabytes of data will be created every day in 2024 with an increasing speed. Under these circumstances, collecting web data and doing traditional human cleaning is no longer feasible.

---

\*equal contribution, zihengq@comp.nus.edu.sg, 20b903054@stu.hit.edu.cn.

†corresponding authors, h.yao@hit.edu.cn, {kai.wang, youy}@comp.nus.edu.sg

Web data usually has complex sources and various qualities, so directly adding them to the training would reduce the benefit of already collected high-quality data, as web data are likely to have noise and redundancy. The problem of how to efficiently utilize web data is worth studying. BLIP [15] adopts bootstrapping to filter and relabel noisy web data based on the model’s self-prediction and human-annotated finetuning. Its remarkable results suggest that their data processing greatly improves data quality. However, BLIP involves model training on all collected data, which may be difficult to scale to even larger data scales.

To scale data processing to an even larger scale, we define the problem in a recursive way. Given a set of data already collected, a new data point can be divided into 3 types: (1) out-of-distribution (noisy) data, (2) redundant data, and (3) clean data with new information. Noisy data need to be filtered out or relabeled; redundant data provide marginal benefit compared to clean data with new information, so its sampling frequency should be correspondingly lowered. Following the above analysis, creating a good dataset can be interpreted as repeatedly adding one informative sample to an already maintained dataset.

Based on this idea, we propose InfoGrowth, a novel framework for efficient online dataset growth with cleanness and diversity awareness. It takes the streaming data as input and returns a clean and diverse dataset. Specifically, for adding a new data point, we estimate its noisy degree and redundancy according to its neighborhood in the already collected dataset. Noisy samples are then relabeled, whereas redundant samples are sampled less frequently or removed. The near-neighbor search is based on an efficient online algorithm, so InfoGrowth has better efficiency and scalability than the previous dataset cleaning algorithms.

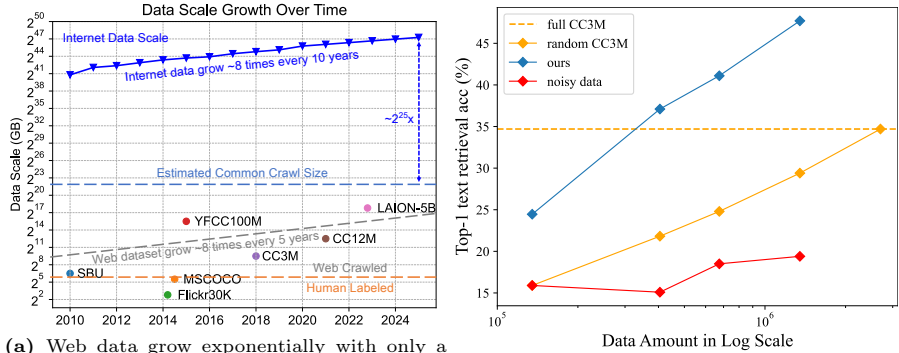
InfoGrowth works for both single-modality tasks and multi-modality tasks across different model architectures. It can achieve 2x~4x data efficiency in both settings, even surpassing its offline counterparts. It is designed to be efficient and scalable, which would be practical for real-world data engines. We hope this work will benefit the dataset preparation of deep learning and sustainable AI.

## 2 Background and Related Works

### 2.1 Deep Learning Datasets

Datasets are a crucial part of deep learning, serving as the foundational bedrock upon which deep learning models are trained and validated. Single-modal datasets emerge since the beginning stage of deep learning. MNIST [13], CIFAR10/100 [12], ImageNet [4], COCO [18] are the most classical datasets for computer vision, which were all human-labeled, with size ranging from ~10k to 10M. Later datasets with much larger sizes like JFT [28] are labeled automatically by algorithm and have a size from 300M to billion, which is far beyond human capability to label. This transition from human labeling to algorithm labeling also happens in other fields like multi-modal datasets.

Multi-modal Datasets encompass diverse data types such as images, text, speech, and audio with paired relations. They can capture rich and nuanced



(a) Web data grow exponentially with only a small fraction used for deep learning. Human labeling/cleaning is no longer feasible for web-scale datasets. (b) Data cleanliness is important to training. Better processing can lead to much higher performance at the same data scale.

**Fig. 1:** Data grows rapidly on an exponential scale. A dataset with cleanliness and diversity leads to better data efficiency and training results.

information. COCO [18], Flickr30K [31], and Visual Genome [11] are human-labeled datasets with high-quality image text data pairs. SBU [20] collects data from the internet and does a cleaning to select useful pairs. YFCC100M [29] filters 100 million media objects from Flickr with certain criteria, with metadata paired. CC3M [26] extract and filter image caption annotations from billions of web pages using Google APIs; with the same procedure and relaxed constraints, CC12M [3] was created. LAION [25] filter data from Common Crawl [1] with CLIP [23] score, resulting in a 400M and 5B dataset respectively. WebVid [2] collects 10 million video clips with captions from the internet. We can see when the dataset size comes to hundreds of millions, human cleaning/labeling is almost infeasible, and automatic methods are employed to craft most extra-large scale datasets. Meanwhile, extra-large scale datasets like LAION-5B remain almost uncurated, let alone Common Crawl [1] which grows in billion pages monthly.

2.2 Web Scale Data Curation

To reduce the influence of noise in web data, researchers have proposed many techniques. For multimodal datasets, the methods mainly include soft-labeling, filtering, and recaptioning. ALBEF [16] utilizes a momentum model to provide pseudo-targets and take weighted mean with the original target, serving as denoised soft labels. BLIP [15] follows the design of ALBEF and further cleans the dataset with captioning and filtering using the same model trained on noisy data and fine-tuned on clean data. It has been shown that much higher data efficiency can be achieved with cleaned data.

For single modal data, DivideMix [17] and Segment Anything [10] are two representative works. DivideMix employs semi-supervised training to fight against noise, training two models together that supervise each other’s data cleanliness and detect the other’s training data noise. In contrast, Segment Anything uses

model-assisted human annotation together with model-labeled data to build a semi-supervised training cycle, aligning the model predictions to human annotations. The majority of these works involve model training on given data and/or human annotation. Segment Anything even train the model for multiple runs.

### 3 How to Make a Dataset Grow

#### 3.1 Preliminaries

Multimodal Models. CLIP [23] is an influential multimodal model that employs natural language supervision for classification tasks. It trains one image encoder and one text encoder with contrastive loss to predict the image-text pairwise relationship, which greatly improves training efficiency compared to traditional supervised learning using labels. Following it, BLIP [15] further adds two other tasks together for training: predicting whether a chosen image-text pair is matching (image-text matching loss), and predicting masked text based on the image (language modeling loss). Another improvement is its soft labeling from ALBEF [16]. These modifications lead to a better capability of BLIP than CLIP, allowing it to perform more vision language tasks (such as captioning, VQA) other than classification or retrieval. More recently, BLIP-V2 [14] utilizes Q-Former to bridge the representation of trained image encoder and large language model. MiniGPT-4 [32] further adds an instruction fine-tuning stage to BLIP-V2 and uses Vicuna [21] as the language model. These modifications equip multimodal models with stronger language skills.

Submodular Function is a set function that describes the decreasing benefit of adding one element to a set. Let  $\Omega$  be a finite set, submodular function  $f$  is defined as:

$$\begin{aligned} f: 2^\Omega \rightarrow \mathbb{R}, \forall x \in \Omega, \forall X, Y \subseteq \Omega \text{ with } X \subseteq Y, \\ f(X \cup \{x\}) - f(X) \geq f(Y \cup \{x\}) - f(Y). \end{aligned} \quad (1)$$

This kind of function is broadly used in many fields. We design a function inspired by it to efficiently estimate the information gain of selected samples.

Online Approximate Near Neighbor Search Algorithm. Traditional submodular function cost  $O(n)$  for each element. This would sum to  $O(n^2)$  overall cost, which is not practical for today’s billion-scale dataset. Therefore to reduce this cost, we introduce online approximate near neighbor search (e.g. HNSW [19]) to do an efficient neighbor search in  $O(\log n)$  time and then calculate a neighborhood-based gain. The submodular property 1 still holds, while neighborhood samples provide a faster estimation of the gain. In even larger scale settings, distributed approximate near neighbor search [5] could be employed.

#### 3.2 Theoretical Analysis

In general, machine learning datasets can be interpreted as conditioned data. Despite there are various types of datasets, it is possible to interpret all datasets

from the same Bayesian inference perspective. For example, classification tasks learn the conditioning of image data on class; multimodal retrieval tasks learn the mutual conditioning of data pair; auto-regressive next token prediction learns the next token distribution conditioned on the previous context. Most datasets are used to capture a conditioned distribution  $P(d|c)$  or  $P(c|d)$  where  $d$  is the data and  $c$  is the interested condition.

For classification tasks, the interested condition  $c$  is the class of given data, and the target is  $P(c|d)$ . Multimodal retrieval tasks are to predict  $P(d_1|d_2)$  and  $P(d_2|d_1)$  where  $d_1, d_2$  are paired data in different modalities; their corresponding condition is each other. Unsupervised datasets are also conditioned on themselves, at either a coarse-grained level (sample-wise, such as contrastive learning) or a fine-grained level (pixel-wise or token-wise, such as MAE and Next Token Prediction). Their intrinsic differences are in the mapping relationships: many-to-one, many-to-many, one-to-many. From this angle of view, datasets of different types can be unified.

Leveraging this insight, we can first map data into an embedding space where the embeddings are separable according to conditions  $c$ , then we can detect the out-of-distribution samples (noise) and redundant samples (low information gain for the task) in the embedding space. A good multimodal model would be a preferred encoder to provide a unified view. Its training objective targets to make the embeddings distributed separably in the space according to natural language. If a multimodal model is strong enough, all data would be theoretically hypersurface separable by human-defined conditions as long as the condition can be defined by natural language.

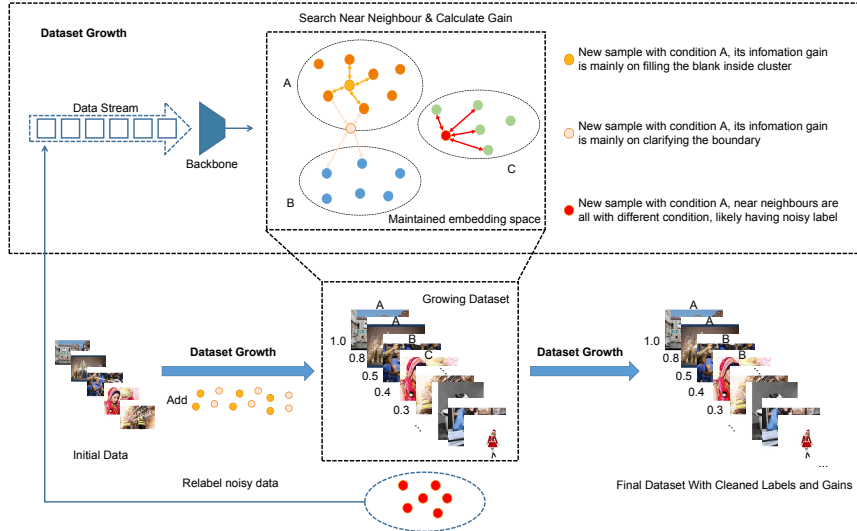
Within an established embedding space, given a set of correctly conditioned data, we can estimate the quality of a new given data point. By adding good data points step by step and efficiently scaling this assessment, it is possible to make datasets grow, with a higher efficiency than “collect then clean”.

### 3.3 Overview of InfoGrowth

Following the theoretical analysis, we design our method InfoGrowth as shown in Fig. 2. Given an input data stream, each data point is first encoded by a backbone model (we use BLIP here) to be an embedding. Then in this embedding space, we design an efficient way to estimate the quality of the data point according to already cleaned data. Clean ones are added to the dataset with calculated "gain", while noisy ones are relabeled and re-estimated. In short, the pipeline mainly consists of: a multimodal encoder model as the backbone, a cleaner for data cleaning, a gain calculator to measure the information gain, and a sampler for choosing samples. The algorithm can be summarized as Alg. 1.

### 3.4 How to detect and correct noisy sample

Cleaner, by its nature, is a model predicting  $P(c|d)$ . When the cleaner prediction has a lower noise than the given data, its prediction could help improve the label quality. Following BLIP [15], on multimodal datasets, the cleaning process mainly uses filter and recaption to improve the data alignment and reduce



**Fig. 2:** Pipeline of InfoGrowth. Streaming data first goes through cleaner, then gain calculator, and finally selector.

noise. We use a trained BLIP model (vit-base-129M) as the encoder and encode image/text into the multimodal representation. The recaptioner we use is MiniGPT-4. Benefiting from the latest development in LLM and multimodal training, it can provide better recaption with a comprehensive understanding of the images. For single modality tasks, a cleaner would be a trained model predicting  $P(c|d)$  so that it can detect suspicious samples and relabel the sample. A difference in our procedure is that we use the cosine similarity of two modalities' encoded features to filter samples whose multimodal cosine similarity is smaller than  $\delta$ . This avoids the overhead of using another encoder in BLIP to identify unmatched samples. The  $\delta$  value can be decided in two ways: use a predefined value or update it with online data statistics.

### 3.5 Gain Calculation

Inspired by submodular functions, we design a submodular-like function for calculating the information gain of a data pair  $(d, c)$ . Traditional submodular functions take  $O(n)$  time to calculate a sample's information gain, which would be a huge overhead for growing billion-scale datasets. To efficiently calculate the gain, we exploit an online approximate nearest neighbor search architecture HNSW [19] to query the nearest  $k$  neighbors of a given data point. The cost of query and update are both  $O(\log(n))$  where  $n$  is the number of collected data points. Then we calculate the cosine distance of those neighbors to the given data, using the mean of distance as the information gain:  $Gain_{Info} =$

---

Algorithm 1 InfoGrowth Algorithm (for 1 step)

---

**Input:** Streaming data pairs  $(d_i, c_i)$ , neighbourhood hyperparameter  $k$ , list  $L$ , NearNeighbour search architecture  $NN$

**Output:** Processed data list  $L$  with gain

```

 $p \leftarrow \text{Cleaner}(c_i, d_i)$ 
if  $p < \delta$  then
   $c_i \leftarrow \text{Cleaner}(d_i)$ 
   $p \leftarrow \text{Cleaner}(c_i, d_i)$ 
  if  $p < \delta$  then
    Continue
  end if
end if
 $d_{\text{feature}} \leftarrow \text{Encoder}(d_i)$ 
 $\text{Neighbors} = NN(d_{\text{feature}}, k)$ 
 $g_i \leftarrow \text{mean}(\text{cosine-distance}(d_{\text{feature}}, \text{Neighbors}))$ 
Add  $(d_i, c_i, g_i)$  to  $L$ , sample based on  $g_i$  values
Add  $d_{\text{feature}}$  to  $NN$ 

```

---

$\text{mean}_i([\text{cos-dis}(d, \text{neighbour}_i)])$ . This design provides a submodular-like property ( $U \subseteq V \Rightarrow f(x, U) \geq f(x, V)$ ) with higher efficiency ( $O(\log n)$  compared to set operation  $O(n)$ ). It also has a good awareness of the neighborhood to avoid redundancy. We then use the gain values for sampling, which further improves the data efficiency beyond cleaning.

On the classification task, we add an "entropy gain" into consideration. It is defined to reflect how much a new sample helps to clarify classification boundary:  $\text{Gain}_{\text{Entropy}} = 1 - p$ , where  $p$  is the probability of using the near neighbor's class label to correctly predict the given data's label. This gain puts more weight on boundary data points than center data points.

### 3.6 How to select informative samples

We design two kinds of samplers: static and dynamic. The static one aims to get better diversity with a fixed total amount of data, while the dynamic one aims to get better computational efficiency on more collected data. Our static sampling sets the probability of choosing a sample as:

$$\mathcal{P}(x_i) = \frac{G_i}{\sum_{j=1}^{|L|} G_j}. \quad (2)$$

Based on the probability given by Eq. (2), we randomly choose samples without replacement to the target number. Our dynamic sampling uses a two-phase sampling strategy, with one phase for overall diversity and another phase for generalization (local diversity). The first phase has the same probability as Eq. (2), and use a total amount of  $\sum_{j=1}^{|L|} G_j$ ; the second phase use the following

probability:

$$G'(x_i) = \max(0.1, 1 - G_i), \mathcal{P}(x_i) = \frac{G'_i}{\sum_{j=1}^{|L|} G'_j}, \quad (3)$$

and sample a total amount of  $\lfloor \sum_{j=1}^{|L|} G'_j \rfloor$ . In the training, the two phases take turns alternately at each epoch. This leads to a cost between 50% and 55% of training with access to all data.

## 4 Experiments

We verify the effectiveness of our method on multiple popular datasets with different architectures and tasks, including 1. BLIP pretraining on multimodal dataset CC3M and retrieval/caption evaluation. 2. ResNet-50 classification on single-modal dataset ImageNet-1K [4]

### 4.1 Datasets and Implementation Details

CC3M is a dataset with web-crawled image-text pairs, filtered with Google API keeping only 0.1% of the raw data. This dataset contains a total of 2.7 million images. We use it to demonstrate the effect of our algorithm in different situations. It is the default dataset for ablation experiments.

ImageNet is a dataset with web-crawled images and human-annotated labels across a large number of classes. We use ImageNet-1K which contains 1000 classes and 1.28M images to evaluate our method on single-modality.

Implementation Details. For backbone encoding, the process is parallelizable without constraint on accelerator number. CC3M cleaner uses MiniGPT-4l which requires at least 1 NVIDIA A100 GPU. Other parts of the pipeline (gain calculation and sampling) can be run on the CPU. For the training of BLIP, we use PyTorch implementation from the official version and train it on 8 NVIDIA A100 GPUs. The model is pre-trained for 20 epochs with a batch size of 640 and an AdamW [27] optimizer with a weight decay of 0.05. We follow the default learning rate warm-up from 1e-6 to 3e-4 and a linear decay with a rate of 0.9. For image data augmentation, we utilize RandAugment [8] except for color inversion. During pre-training, we randomly crop images to a resolution of  $224 \times 224$ . Other implementation details can be found in the supplementary material.

### 4.2 Evaluation on Multimodal Dataset

Image-text retrieval consists of two parts: (1) image as the query to retrieve texts (TR); (2) text as the query to retrieve images (IR). The pre-trained model is evaluated on MSCOCO [18] and Flickr30K [31]. For the COCO zero-shot setting, the pre-trained model is directly evaluated on the test data. For zero-shot retrieval on Flickr30K, we follow the procedure proposed in [15] (zero-shot

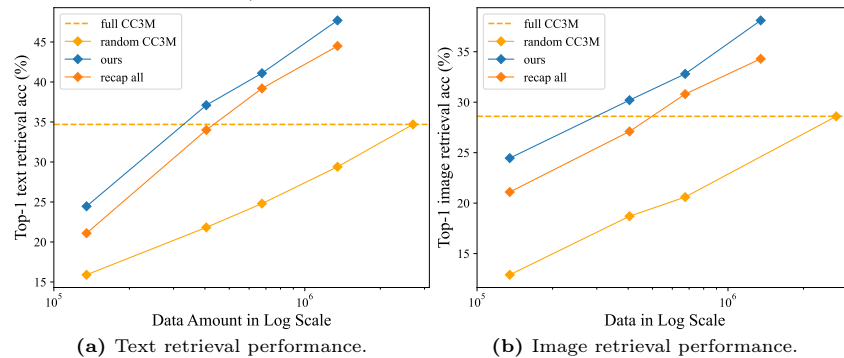


**Table 1: Fine-tuning and zero-shot image-text retrieval results on MSCOCO and Flickr30K.** InfoGrowth demonstrates an improved data efficiency.

Model	Dataset Size	MSCOCO (5K test set)						Flickr30K (1K test set)						
		Image→Text			Text→Image			Image→Text			Text→Image			
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	
BLIP	CC3M	2.71M	70.9	90.6	95.3	54.3	79.8	87.4	89.2	98.4	99.6	75.1	93.2	96.4
	Ours	0.40M	<b>70.9</b>	<b>91.3</b>	<b>95.9</b>	53.3	79.5	87.0	84.6	97.1	98.7	69.5	90.3	94.5
	Ours	0.68M	<b>71.5</b>	<b>91.1</b>	<b>95.3</b>	<b>53.9</b>	<b>79.5</b>	<b>87.5</b>	89.0	98.4	99.0	74.6	92.9	95.9
	Ours	1.35M	<b>72.1</b>	<b>92.2</b>	<b>96.1</b>	<b>55.3</b>	<b>81.4</b>	<b>89.0</b>	<b>91.5</b>	<b>99.1</b>	<b>99.8</b>	<b>77.7</b>	<b>94.0</b>	<b>96.9</b>
	CC3M	2.71M	34.7	60.7	71.9	28.0	52.5	63.7	87.4	98.2	99.5	72.8	90.7	94.6
	Ours	0.40M	<b>37.1</b>	<b>63.8</b>	<b>74.2</b>	<b>28.1</b>	<b>53.8</b>	<b>65.1</b>	82.2	96.9	98.6	67.0	87.8	92.7
	Ours	0.68M	<b>41.1</b>	<b>67.7</b>	<b>77.5</b>	<b>32.8</b>	<b>58.2</b>	<b>68.8</b>	86.8	97.7	99.0	70.2	90.2	90.4
Ours	1.35M	<b>47.7</b>	<b>73.0</b>	<b>82.0</b>	<b>38.1</b>	<b>63.3</b>	<b>73.2</b>	<b>88.7</b>	<b>98.3</b>	99.0	<b>74.2</b>	<b>92.2</b>	<b>95.4</b>	
CLIP	CC3M	2.71M	60.4	85.3	93.2	<b>48.9</b>	75.4	84.7	77.3	91.1	<b>93.2</b>	71.6	90.1	91.4
	Ours	0.68M	<b>62.2</b>	<b>86.4</b>	<b>94.0</b>	48.6	<b>76.3</b>	<b>86.0</b>	<b>80.5</b>	<b>94.6</b>	<b>98.3</b>	<b>72.3</b>	<b>90.6</b>	<b>93.0</b>

evaluation on Flickr with the model fine-tuned using MSCOCO). For the fine-tuning setting, the model is further fine-tuned on the Flickr30K training data and evaluated on the validation/test data.

As seen in Fig. 3, we train the BLIP model on the dataset processed by our method and the original CC3M. Our processed dataset uses less data and has a much higher data efficiency than using the original CC3M. With only 0.4M data (0.8M seen), our static dataset can surpass the training performance using the full original CC3M. Moreover, with our gain-based dynamic two-phase sampling (saving at least 45% training cost), the training takes only 8.2% of the CC3M’s computation and achieves a 3 percent zero-shot text retrieval performance improvement on COCO. When the sample number increases, this gap is further enlarged. Compared to only recaptioning samples, InfoGrowth shows higher data efficiency with gain-based sampling. In Table 1, we can see the dataset processed by InfoGrowth demonstrates a strong zero-shot transfer ability on MSCOCO, using less than 1/6 data to surpass original CC3M. On Flickr30K, the data needed is increased to 1/2.

**Fig. 3:** InfoGrowth demonstrates better training results than using original CC3M and MiniGPT4 recaptioned CC3M at the same amount of data. Its cost is also much lower than recaptioning all data.

**Downstream Tasks** In the next paragraphs, we evaluate our trained model on downstream cross-modal tasks. We first introduce these tasks:

**Table 2: Comparison with BLIP model pre-trained on different data sources for VQA, NLVR<sup>2</sup>, RefCOCO+ and COCO Captioning.** CLIP architecture can not be evaluated on part of these tasks due to structural limitations.

Dataset	#Samples	VQA		NLVR <sup>2</sup>		COCO Caption		Nocaps Validation	
		test-dev	test-std	dev	test-P	B@4	CIDEr	CIDEr	SPICE
Random- <i>CC3M</i>	0.67M	68.3	66.2	73.6	73.8	35.9	118.1	85.3	11.6
<i>CC3M</i> [26]	2.7M	71.3	71.4	75.9	76.5	36.1	121.4	90.6	12.8
Ours	0.67M	70.9 <sub>-0.4</sub>	71.0 <sub>-0.4</sub>	75.4 <sub>-0.5</sub>	76.8 <sub>+0.3</sub>	36.7 <sub>+0.6</sub>	121.7 <sub>+0.3</sub>	89.8 <sub>-0.8</sub>	13.1 <sub>+0.3</sub>
Ours	1.35M	71.6 <sub>+0.3</sub>	71.9 <sub>+0.5</sub>	76.1 <sub>+0.2</sub>	77.4 <sub>+0.9</sub>	36.9 <sub>+0.8</sub>	122.3 <sub>+0.9</sub>	90.6 <sub>+0.0</sub>	13.3 <sub>+0.5</sub>

Visual Question Answering (VQA.) We evaluate our model’s performance on the VQA [7] task, where the model needs to provide an answer based on an image and a question. We consider it as an answer generation task that allows open-vocabulary VQA for better results, following previous works [22, 23]. The results are presented in Table 2. The BLIP trained on InfoGrowth outperforms baseline by InfoGrowth% on test-dev splits, demonstrating the effectiveness of our InfoGrowth for improving VQA performance.

Visual Reasoning. The Natural Language Visual Reasoning task is a binary classification task that requires the model to reason about two images and a question in natural language. The dataset (*NLVR*<sup>2</sup>) [27] contains 107,292 examples of human-written English sentences paired with web photographs. Multi-modal reasoning is crucial for the completion of this task. We observe that BLIP trained on our dataset achieved InfoGrowth% accuracy compared to InfoGrowth% achieved by the CC3M, as shown in Table 2.

Image captioning. The task involves describing an input image, which we evaluate using No-Caps and COCO datasets. Both datasets are fine-tuned on COCO with the Language Modeling (LM) loss. We adopt a zero-shot setting for No-Caps dataset, and start each caption with the phrase “a picture of” for the BLIP architecture. We do not pre-train using COCO to avoid information leakage. Our results outperform the baseline with a much smaller quantity of pre-training data, as shown in Table 2.

### 4.3 Evaluation on Single-modal Dataset

Besides the multimodal datasets, InfoGrowth can also be applied to single-modal datasets as analyzed in Sec. 3.2. A distribution of  $P(d|c)$  can be established in the embedding space, and the target objective  $P(c|d)$  can be transferred by Bayesian’s rule to  $P(d|c)P(c)/P(d)$ , which is proportional to  $P(d|c)$ . Therefore, the corresponding data quality can be estimated within the same framework as in the previous subsection.

ImageNet Classification. We evaluate InfoGrowth on ImageNet-1k [4] Classification task using ResNet-50 [8]. The corresponding gain is defined in Sec. 3.5 as the mean of information gain and entropy gain. The corresponding result is shown in Table 3. One can find that InfoGrowth not only significantly surpasses traditional static dataset compression methods like GC and EL2N in both performance and computation cost, but also outperform dynamic pruning methods like UCB in final accuracy. These results are suggesting InfoGrowth’s generality and competitiveness across task, data, and model architecture.

**Table 3:** Comparison of performance and time cost on ImageNet-1K. Results are reported with ResNet-50 under 50% prune ratio for 90 epochs on an 8-A100-GPU server. “Time” is wall clock time; “Total (n\*h)” is the total node hour.

	GC [9]	EL2N-20 [30]	UCB [24]	InfoGrowth	Full Data
Acc (%)	72.8±0.4	74.6±0.4	75.3±0.3	<b>75.8±0.3</b>	76.4±0.2
Training (h)	8.75	8.75	8.75	8.75	17.5
Overhead (h)	>24	>17.5	0.03	0.8	0.0
Total (n*h)	>94	>210	<b>70</b>	<b>70.8</b>	140.0

#### 4.4 Ablation

If not stated otherwise, we train BLIP with 0.4M CC3M data processed with the specified ablation setting and all other settings as default, and evaluate on the COCO zero-shot retrieval task.

Evaluating the components of InfoGrowth. We design an ablation study to investigate the characteristics of cleaner and sampler modules in table 4. We train BLIP on the dataset processed by InfoGrowth with different components ablated and evaluate its zero-shot performance on the COCO retrieval task. The baseline is a random selection from CC3M. We can see each module provides a performance gain compared to the baseline, and using caption+filtering cleaning together with gain-based sampling achieves the best performance in our experiment. With only 0.4M data ever seen, BLIP trained on InfoGrowth’s processed dataset can achieve comparable retrieval performance with using the full CC3M. Using dynamic sampling, the training cost can be further saved by nearly half, taking only 8% of the cost on full CC3M.

**Table 4: Component ablation.** Both the cleaning and gain-based sampling operations are important to the downstream retrieval. With only 0.4M data seen, InfoGrowth is able to beat full CC3M(2.7M). **Table 5: k neighbor and Mean method.** Comparison of neighborhood size and arithmetic/harmonic mean to calculate gain. k=4 with mean performing the best in the experiment.

	recaption	sampling	time(h)	TR@IR@1		arithmetic		harmonic		
				TR@1	IR@1	TR@1	IR@1	TR@1	IR@1	
			4.2	21.8	18.7	k	TR@1	IR@1	TR@1	IR@1
✓			4.2	34.1	27.1	k=1	33.5	26.7	33.5	26.7
	✓		2.3	23.9	21.0	k=2	33.5	27.2	34.6	27.2
✓	✓		2.3	<b>37.1</b>	<b>28.1</b>	k=4	<b>37.1</b>	<b>28.1</b>	35.3	27.7
		Full CC3M	14	34.7	28.0	k=8	33.6	26.7	33.2	27.1

Influence of neighborhood query number k and way of calculating mean. In the gain calculator, we use the neighborhood of a given feature to calculate its information gain. Here we conduct an ablation experiment in Tab. 5 to show the influence of different k choices. We train the BLIP model with the dataset processed by different k values. We can see that when using the arithmetic mean, increasing k from 1 to 4 leads to a gradual performance improvement. Further increasing k to 8 shows a performance degradation. When

using harmonic mean, more weight is put on closer neighbors, and increasing  $k$  gradually improves the performance, as it approximates the submodular gain on full data with higher precision.

**Table 6: Gain Function.** Averaged image and text gain function is best for overall performance. **Table 7: Noise ratio.** Our method provides noise resistance.

case	TR@1 IR@1		noise ratio	method	TR@IR@1	
	TR@1	IR@1	0%	clean baseline	34.7	28.0
Image Gain	32.8	28.0	10%	unprocessed	17.8	15.2
Text Gain	34.0	26.7	10%	ours	33.8	26.6
Info+Alignment Gain	33.7	27.2	25%	unprocessed	15.0	13.1
Info Gain	<b>37.1</b>	<b>28.1</b>	25%	ours	31.2	25.6

Multimodal data gain calculation. The default multimodal gain (Info Gain) calculation has two parts: image gain and text gain, each calculated with the mean of  $k$  nearest neighbors’ cosine distance. Alignment Gain is the cosine similarity between two modalities’ embedding. We ablate each part in Tab. 6 and see their performance difference. We can see that with gain from both modalities, the dataset processed gives the best training result on the retrieval task. Adding Alignment Gain into gain calculation doesn’t improve the retrieval performance.

Data amount. In general, more collected data leads to better training results. From Fig. 1b and 3, three qualities of data from InfoGrowth, recaption, and original each form a linear line in exponential data growth. However there are some interesting findings to be discussed in Sec. 5.3.

Model generalization and Task generalization. In Table 1, we test the performance of InfoGrowth on CLIP architecture. It suggests the effectiveness of growing a dataset within the embedding space of multimodal models is independent to model architecture.

In Table 8, we can see that on MSRVT, which is video classification with different data distribution to BLIP training data, our trained BLIP on 0.67M CC3M data processed by InfoGrowth shows improved ability than trained on the original CC3M. InfoGrowth also improves over the original CC3M on ImageNet-1k zero-shot classification with 1.35M processed data. It suggests that task generalizability is inherited by InfoGrowth processed dataset.

**Table 8:** Video classification on MSRVT and Image classification on ImageNet-1K for BLIP.

Dataset	Scale	MSRVT			Imagenet
		R@1↑	R@5↑	R@10↑	Acc↑
Rand-cc3m	0.67M	23.3	42.8	53.3	57.3
cc3m	2.82M	26.0	46.3	58.0	62.5
Ours	0.67M	26.4	47.3	58.4	62.1
Ours	1.35M	<b>28.8</b>	<b>49.4</b>	<b>60.0</b>	<b>62.9</b>

Data noisy scale. InfoGrowth demonstrates the ability to resist noise as shown in table 7. When more than 25% of the captions are randomly shuffled, training directly on it shows a degraded performance. In contrast, InfoGrowth is still able to denoise the data and achieve a reasonable.

## 5 Discussion

### 5.1 Comprehensive Comparison with Other Methods

**Table 9:** A comprehensive comparison with other methods.

method	type	functionality	modality	per sample cost
InfoGrowth	online	clean and select	both	Inference, $O(\log n)$ growth
Bootstrapping	offline	clean	multi-modal	Training*2+Inference
Data Pruning	offline	select	single-modal	Inference, $O(n)$ growth

InfoGrowth mainly differs from existing dataset processing methods in several aspects: 1. it is an online algorithm 2. it supports data for both single-modal and multimodal tasks 3. it is efficiently scalable to support the data growth.

The cost of InfoGrowth is analyzed as follows:  $O(M_{encoder})$  per sample to make the decision,  $O(kd \cdot \log n)$  for search and compare  $k$  neighbors with embedding size  $d$ ,  $O(M_{cleaner})$  to clean detected noisy samples ( $\sim 30\%$  samples in CC3M).

We compare the difference of InfoGrowth with other methods in Table 9. When taking into consideration that data grow on an exponential scale, bootstrapping methods like BLIP could be still acceptable because they produce a model; data pruning methods with  $O(n)$  per sample cost would take  $O(n^2)$  on an exponentially growing  $n$ , and this cost occurs periodically. In contrast, InfoGrowth doesn't incur further costs on already collected data, and deals with the new data in an  $O(\log n)$  growth.


A most recent work InfoBatch [22] proposes to do unbiased dynamic data pruning to losslessly accelerate deep learning training with negligible cost. It would introduce little cost when more data is added. However, InfoBatch currently cannot generalize to multimodal tasks because its rescaling doesn't work for contrastive loss, while most multimodal tasks use contrastive learning.

### 5.2 Scaling to Even Larger Scale

Currently, the experiments verify InfoGrowth on a million scale of data using HNSW. To further scale to billion-scale data and real-world data pipeline, InfoGrowth only needs a little modification. Note that data encoding and cleaner prediction could be totally asynchronous. Only the near-neighbor search architecture needs to be synchronized, so a distributed near-neighbor search algorithm [5] would be enough for further scaling.

### 5.3 The Marginal Benefit of Data

From Fig. 1b and 3, one can see that when data of the same quality is added on an exponential scale (computation increased proportionally), the model performance increases almost linearly. Sun et al. [28] reported the same phenomenon which motivated



**Fig. 4:** Gain decay with more data collected.

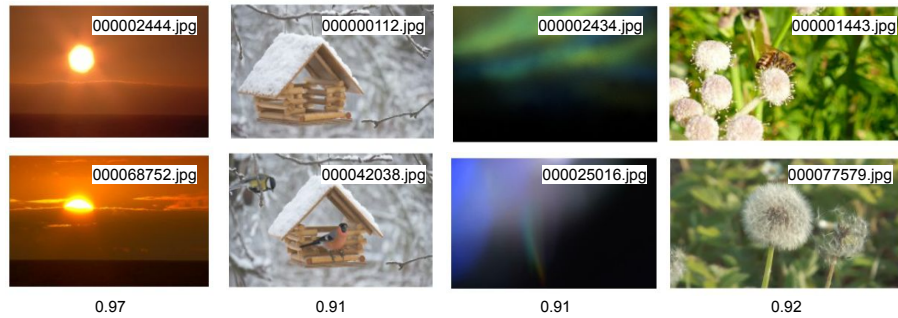
them to create the ultra-large-scale dataset JFT. This also implies that with more data collected, the further gain of added data would decrease. Our defined gain also reflects this trend, which is demonstrated in Fig. 4.

Another interesting phenomenon of InfoGrowth is that when the total selected data amount is fixed, choosing from more candidate data doesn't improve performance when enough data is seen (Table. 10). This implies that the data quality improvement of InfoGrowth has an upper bound because marginal decay.

**Table 10:** InfoGrowth with 0.67M selected data.

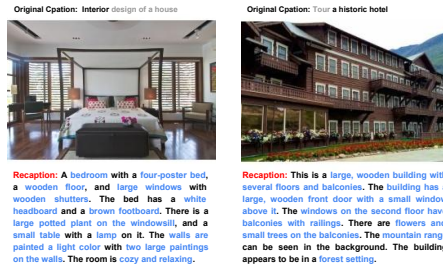
Seen Data	TR@1	IR@1
1.35M	41.1	32.8
2.7M	40.9	32.5

#### 5.4 Visualizations



**Fig. 5:** Redundant samples detected in CC3M and their cosine similarity.

In this section, we visualize in Fig. 5 and 6 some examples from our algorithm's detected redundant/noisy samples. The static selection according to gain and dynamic two-phase sampling both prevent those redundant samples from occurring at the same time. For noisy samples, the cleaner will improve the data quality with recaptioning/relabeling.



**Fig. 6:** Cleaner improving data quality. Grey words are low-quality parts and blue words are high-quality descriptions.

## 6 Conclusion

In this work, we propose a novel solution to efficiently deal with the growing data demand of deep learning. By turning the data cleaning and selection process into an online algorithm, our proposed InfoGrowth algorithm demonstrates its competitive capability in improving data quality for both multi-modal tasks and single-modal tasks, even competitive with offline counterparts. The framework design of InfoGrowth allows it to be scaled to an even larger data scale. We hope this new kind of online efficient algorithm will contribute to the creation of high-quality large-scale deep learning datasets and benefit sustainable AI.

## Acknowledgements

This work is supported by Alibaba Group through Alibaba Innovative Research Program. This work is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-PhD-2021-08-008). Yang You’s research group is being sponsored by NUS startup grant (Presidential Young Professorship), Singapore MOE Tier-1 grant, ByteDance grant, ARCTIC grant, SMI grant (WBS number: A-8001104-00-00), Alibaba grant, and Google grant for TPU usage. This work is partially supported by the National Natural Science Foundation of China (62176165), the Stable Support Projects for Shenzhen Higher Education Institutions (20220718110918001), the Natural Science Foundation of Top Talent of SZTU(GDRC202131). This work was supported in part by the National Science and Technology Major Project (2021ZD0110901). We thank Haonan Wang for valuable discussions and feedback.

## References

1. Common crawl, <https://commoncrawl.org/>
2. Bain, M., Nagrani, A., Varol, G., Zisserman, A.: Frozen in time: A joint video and image encoder for end-to-end retrieval. In: IEEE International Conference on Computer Vision (2021)
3. Changpinyo, S., Sharma, P., Ding, N., Soricut, R.: Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. CoRR **abs/2102.08981** (2021), <https://arxiv.org/abs/2102.08981>
4. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 248–255 (2009). <https://doi.org/10.1109/CVPR.2009.5206848>
5. Douze, M., Guzhva, A., Deng, C., Johnson, J., Szilvasy, G., Mazaré, P.E., Lomeli, M., Hosseini, L., Jégou, H.: The faiss library (2024)
6. Duarte, F.: Amount of data created daily (2024) (Dec 2023), <https://explodringtopics.com/blog/data-generated-per-day>
7. Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., Parikh, D.: Making the v in vqa matter: Elevating the role of image understanding in visual question answering (2017)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. CoRR **abs/1512.03385** (2015), <http://arxiv.org/abs/1512.03385>
9. Iyer, R., Khargoankar, N., Bilmes, J., Asanani, H.: Submodular combinatorial information measures with applications in machine learning. In: Algorithmic Learning Theory. pp. 722–754. PMLR (2021)
10. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. arXiv:2304.02643 (2023)
11. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., Bernstein, M.S., Li, F.F.: Visual genome: Connecting language and vision using crowdsourced dense image annotations (2016)

12. Krizhevsky, A., Nair, V., Hinton, G.: Cifar-10 (canadian institute for advanced research) <http://www.cs.toronto.edu/~kri/cifar.html>
13. LeCun, Y., Cortes, C.: MNIST handwritten digit database (2010), <http://yann.lecun.com/exdb/mnist/>
14. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models (2023)
15. Li, J., Li, D., Xiong, C., Hoi, S.: BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., Sabato, S. (eds.) Proceedings of the 39th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 162, pp. 12888–12900. PMLR (17–23 Jul 2022), <https://proceedings.mlr.press/v162/li22n.html>
16. Li, J., Selvaraju, R.R., Gotmare, A.D., Joty, S.R., Xiong, C., Hoi, S.C.H.: Align before fuse: Vision and language representation learning with momentum distillation. CoRR **abs/2107.07651** (2021), <https://arxiv.org/abs/2107.07651>
17. Li, J., Socher, R., Hoi, S.C.H.: Dividemix: Learning with noisy labels as semi-supervised learning (2020)
18. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13. pp. 740–755. Springer (2014)
19. Malkov, Y.A., Yashunin, D.A.: Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. CoRR **abs/1603.09320** (2016), <http://arxiv.org/abs/1603.09320>
20. Ordonez, V., Kulkarni, G., Berg, T.: Im2text: Describing images using 1 million captioned photographs. In: Shawe-Taylor, J., Zemel, R., Bartlett, P., Pereira, F., Weinberger, K. (eds.) Advances in Neural Information Processing Systems. vol. 24. Curran Associates, Inc. (2011), [https://proceedings.neurips.cc/paper\\_files/paper/2011/file/5dd9db5e033da9c6fb5ba83c7a7e99-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2011/file/5dd9db5e033da9c6fb5ba83c7a7e99-Paper.pdf)
21. Peng, B., Li, C., He, P., Galley, M., Gao, J.: Instruction tuning with gpt-4 (2023)
22. Qin, Z., Wang, K., Zheng, Z., Gu, J., Peng, X., Xu Zhao Pan, Zhou, D., Shang, L., Sun, B., Xie, X., You, Y.: Infobatch: Lossless training speed up by unbiased dynamic data pruning. In: The Twelfth International Conference on Learning Representations (2024), <https://openreview.net/forum?id=C61sk5Lsk6>
23. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision (2021)
24. Raju, R.S., Daruwalla, K., Lipasti, M.: Accelerating deep learning with dynamic data pruning (2021). <https://doi.org/10.48550/ARXIV.2111.12621>, <https://arxiv.org/abs/2111.12621>
25. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., Schramowski, P., Kundurthy, S., Crowson, K., Schmidt, L., Kaczmarczyk, R., Jitsev, J.: Laion-5b: An open large-scale dataset for training next generation image-text models (2022)
26. Sharma, P., Ding, N., Goodman, S., Soricut, R.: Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In: Gurevych, I., Miyao, Y. (eds.) Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 2556–2565. Association for Computational Linguistics, Melbourne, Australia (Jul 2018). <https://doi.org/10.18653/v1/P18-1238>, <https://aclanthology.org/P18-1238>



27. Suhr, A., Zhou, S., Zhang, A., Zhang, I., Bai, H., Artzi, Y.: A corpus for reasoning about natural language grounded in photographs (2019)
28. Sun, C., Shrivastava, A., Singh, S., Gupta, A.: Revisiting unreasonable effectiveness of data in deep learning era. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 843–852 (2017). <https://doi.org/10.1109/ICCV.2017.97>
29. Thomee, B., Shamma, D.A., Friedland, G., Elizalde, B., Ni, K., Poland, D., Borth, D., Li, L.J.: Yfcc100m: the new data in multimedia research. *Communications of the ACM* **59**(2), 64–73 (Jan 2016). <https://doi.org/10.1145/2812802>, <http://dx.doi.org/10.1145/2812802>
30. Toneva, M., Sordoni, A., Combes, R.T.d., Trischler, A., Bengio, Y., Gordon, G.J.: An empirical study of example forgetting during deep neural network learning (2018). <https://doi.org/10.48550/ARXIV.1812.05159>, <https://arxiv.org/abs/1812.05159>
31. Young, P., Lai, A., Hodosh, M., Hockenmaier, J.: 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). [https://doi.org/10.1162/tacl\\_a\\_00166](https://doi.org/10.1162/tacl_a_00166), <https://aclanthology.org/Q14-1006>
32. Zhu, D., Chen, J., Shen, X., Li, X., Elhoseiny, M.: Minigpt-4: Enhancing vision-language understanding with advanced large language models (2023)